ONTAP 9.3 commands

ONTAP commands

NetApp
February 14, 2025

This PDF was generated from https://docs.netapp.com/us-en/ontap-cli-93/index.html on February 14, 2025. Always check docs.netapp.com for the latest.

Table of Contents

ONTAP 9.3 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 . . . 8
- Related Links . . . 10

top . . . 11
- Description . . . 11
- Examples . . . 11

up . . . 12
- Description . . . 12
- Examples . . . 12

application commands . . . 13
- application create . . . 13
- application delete . . . 14
- application modify . . . 14
- application show-detail . . . 16
- application show-statistics . . . 24
- application show-template . . . 27
- application show . . . 28
- application provisioning commands . . . 29
- application snapshot commands . . . 30

application volume commands . . . 34
autobalance commands . . . 37
autobalance aggregate commands . . . 37
autobalance volume commands . . . 45
cluster commands . . . 49
cluster add-node . . . 49
cluster create . . . 50
cluster join . . . 51
cluster modify . . . 52
cluster ping-cluster . . . 53
cluster setup . . . 54
cluster show . . . 59
cluster unjoin . . . 61
cluster contact-info commands . . . 61
cluster date commands . . . 64
cluster ha commands . . . 68
cluster identity commands . . . 69
cluster image commands . . . 70
cluster kernel-service commands . . . 89
cluster log-forwarding commands . . . 91
cluster peer commands . . . 95
cluster quorum-service commands . . . 128
cluster ring commands . . . 130
cluster statistics commands . . . 131
cluster time-service commands . . . 133
event commands . . . 142
event catalog commands . . . 142
event config commands . . . 145
event destination commands . . . 149
event filter commands . . . 156
event log commands . . . 176
event mailhistory commands . . . 180
event notification commands . . . 182
event route commands . . . 196
event snmphistory commands . . . 202
event status commands . . . 204
job commands . . . 208
job delete . . . 208
job pause . . . 208
job resume . . . 209
job show-bynode . . . 210
job show-cluster . . . 211
job show-completed . . . 214
job show . . . 216
job stop . . . 219

job unclaim 220
job watch-progress 220
job history commands 221
job initstate commands 224
job private commands 226
job schedule commands 235
lun commands 247
lun create 247
lun delete 250
lun maxsize 251
lun modify 254
lun move-in-volume 258
lun resize 259
lun show 261
lun bind commands 268
lun copy commands 271
lun igroup commands 278
lun import commands 288
lun mapping commands 299
lun move commands 308
lun persistent-reservation commands 314
lun portset commands 317
lun transition commands 322
metrocluster commands 328
metrocluster configure 328
metrocluster heal 332
metrocluster modify 332
metrocluster show 333
metrocluster switchback 336
metrocluster switchover 337
metrocluster check commands 338
metrocluster config-replication commands 364
metrocluster configuration-settings commands 366
metrocluster interconnect commands 395
metrocluster node commands 405
metrocluster operation commands 410
metrocluster vserver commands 413
network commands 417
network ping 417
network ping6 418
network test-path 420
network traceroute 421
network traceroute6 423
network arp commands 424
network cloud commands 428

network connections commands 429
network device-discovery commands 440
network fcp commands 443
network interface commands 453
network ipspace commands 482
network ndp commands 485
network options commands 493
network port commands 504
network qos-marking commands 525
network route commands 527
network routing-groups commands 533
network subnet commands 538
network tcpdump commands 544
network test-link commands 548
network tuning commands 552
protection-type commands 558
protection-type show 558
qos commands 560
qos adaptive-policy-group commands 560
qos policy-group commands 565
qos settings commands 571
qos statistics commands 573
qos workload commands 612
security commands 615
security snmpusers 615
security audit commands 616
security certificate commands 618
security config commands 637
security key-manager commands 645
security login commands 669
security protocol commands 708
security saml-sp commands 711
security session commands 716
security ssh commands 756
security ssl commands 762
snaplock commands 768
snaplock compliance-clock commands 768
snaplock event-retention commands 769
snaplock legal-hold commands 776
snaplock log commands 783
snapmirror commands 789
snapmirror abort 789
snapmirror break 791
snapmirror create 794
snapmirror delete 801

snapmirror initialize-ls-set . . . 804
snapmirror initialize . . . 806
snapmirror list-destinations . . . 811
snapmirror modify . . . 816
snapmirror promote . . . 820
snapmirror protect . . . 822
snapmirror quiesce . . . 823
snapmirror release . . . 826
snapmirror restore . . . 828
snapmirror resume . . . 835
snapmirror resync . . . 838
snapmirror set-options . . . 844
snapmirror show-history . . . 845
snapmirror show . . . 851
snapmirror update-ls-set . . . 874
snapmirror update . . . 876
snapmirror config-replication commands . . . 881
snapmirror policy commands . . . 887
snapmirror snapshot-owner commands . . . 899
statistics commands . . . 903
statistics show-periodic . . . 903
statistics show . . . 906
statistics start . . . 909
statistics stop . . . 911
statistics aggregate commands . . . 912
statistics cache commands . . . 913
statistics catalog commands . . . 914
statistics disk commands . . . 919
statistics lif commands . . . 921
statistics lun commands . . . 922
statistics nfs commands . . . 923
statistics node commands . . . 950
statistics oncrpc commands . . . 951
statistics port commands . . . 953
statistics preset commands . . . 955
statistics samples commands . . . 962
statistics settings commands . . . 963
statistics system commands . . . 964
statistics top commands . . . 966
statistics volume commands . . . 968
statistics vserver commands . . . 969
statistics workload commands . . . 971
statistics-v1 commands . . . 973
statistics-v1 nfs commands . . . 973
statistics-v1 protocol-request-size commands . . . 1000

storage-service commands . . . 1003
storage-service show . . . 1003
storage commands . . . 1005
storage aggregate commands . . . 1005
storage array commands . . . 1094
storage automated-working-set-analyzer commands . . . 1110
storage bridge commands . . . 1116
storage disk commands . . . 1145
storage encryption commands . . . 1196
storage errors commands . . . 1209
storage failover commands . . . 1211
storage firmware commands . . . 1252
storage iscsi-initiator commands . . . 1259
storage load commands . . . 1264
storage path commands . . . 1267
storage pool commands . . . 1277
storage port commands . . . 1290
storage raid-options commands . . . 1297
storage shelf commands . . . 1298
storage switch commands . . . 1345
storage tape commands . . . 1379
system commands . . . 1409
system chassis commands . . . 1409
system cluster-switch commands . . . 1412
system configuration commands . . . 1422
system controller commands . . . 1435
system feature-usage commands . . . 1485
system fru-check commands . . . 1488
system ha commands . . . 1489
system health commands . . . 1514
system license commands . . . 1530
system node commands . . . 1542
system script commands . . . 1658
system service-processor commands . . . 1663
system services commands . . . 1692
system smtape commands . . . 1732
system snmp commands . . . 1742
system status commands . . . 1752
system timeout commands . . . 1755
template commands . . . 1757
template copy . . . 1757
template delete . . . 1757
template download . . . 1758
template provision . . . 1758
template rename . . . 1760

template show-permissions 1760
template show 1762
template upload 1763
template parameter commands 1764
volume commands 1768
volume autosize 1768
volume create 1769
volume delete 1779
volume expand 1780
volume make-vsroot 1783
volume modify 1784
volume mount 1794
volume offline 1795
volume online 1796
volume rehost 1797
volume rename 1798
volume restrict 1799
volume show-footprint 1799
volume show-space 1807
volume show 1813
volume size 1833
volume transition-prepare-to-downgrade 1834
volume unmount 1835
volume aggregate commands 1835
volume clone commands 1836
volume efficiency commands 1853
volume encryption commands 1877
volume file commands 1883
volume flexgroup commands 1912
volume inode-upgrade commands 1914
volume move commands 1917
volume qtree commands 1933
volume quota commands 1945
volume reallocation commands 1978
volume schedule-style commands 1987
volume snaplock commands 1987
volume snapshot commands 1993
volume transition-convert-dir commands 2025
vserver commands 2028
vserver add-aggregates 2028
vserver add-protocols 2028
vserver context 2029
vserver create 2030
vserver delete 2032
vserver modify 2033

vserver prepare-for-revert . . . 2035
vserver remove-aggregates . . . 2035
vserver remove-protocols . . . 2036
vserver rename . . . 2036
vserver restamp-msid . . . 2037
vserver show-aggregates . . . 2038
vserver show-protocols . . . 2039
vserver show . . . 2040
vserver start . . . 2044
vserver stop . . . 2045
vserver unlock . . . 2046
vserver active-directory commands . . . 2046
vserver audit commands . . . 2052
vserver check commands . . . 2062
vserver cifs commands . . . 2069
vserver config-replication commands . . . 2217
vserver data-policy commands . . . 2219
vserver export-policy commands . . . 2223
vserver fcp commands . . . 2266
vserver fpolicy commands . . . 2285
vserver group-mapping commands . . . 2343
vserver iscsi commands . . . 2350
vserver locks commands . . . 2392
vserver migrate commands . . . 2402
vserver name-mapping commands . . . 2413
vserver nfs commands . . . 2421
vserver peer commands . . . 2465
vserver san commands . . . 2483
vserver security commands . . . 2484
vserver services commands . . . 2550
vserver smtape commands . . . 2712
vserver snapdiff-rpc-server commands . . . 2713
vserver vscan commands . . . 2715
Legal notices . . . 2784
Copyright . . . 2784
Trademarks . . . 2784
Patents . . . 2784
Privacy policy . . . 2784
Open source . . . 2784

ONTAP 9.3 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
- `MB` - Megabytes
- `GB` - Gigabytes
- `TB` - Terabytes
- `PB` - Petabytes

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.

`[-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 create

(DEPRECATED)-Create an application

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 of Data ONTAP.

This command creates an application from an application template. It provisions storage for the application and optionally performs additional configuration.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

The Vserver in which the application is created.

`-application <text>` - Application (privilege: advanced)

The name for the application to create.

`-application-template <text>` - Application Template (privilege: advanced)

The application template upon which to base creation of the application.

`[-parameters <key-value-pair>,…]` - Application Template Customization Parameters (privilege: advanced)

A comma separated list of parameter names and their values to customize the application template for the desired application. The specific parameters which may be used to customize an application vary depending on the specified application template.

For each template, there are both required and optional parameters. Attempting to execute the command without a required parameter will result in failure and a message identifying the missing parameter. Command line tab-completion provides lists of parameters and descriptions of the syntax required to specify the parameter values.

A parameter and its value is specified as: *`parameter-name:parameter-value`* . Multiple parameters are specified by separating them with a comma as: *`parameter1-name:parameter1-value,parameter2-name:parameter2-value`* . The correct syntax for specifying a parameter value depends on the type of the parameter.

Some parameters must be specified for multiple instances of the same type of object. These are specified using a zero-based array index to identify each unique object. These types of objects have sub-parameters that provide values for specific attributes of each object. A sub-parameter and its value for the Nth instance of an object is specified as: *`object-type[n].sub-parameter-name:sub-parameter-value`* .

`[-foreground {true|false}]` - Execute in the Foreground (privilege: advanced)

A boolean value that determines if the application creation executes synchronously at the command prompt in the foreground or asynchronously in the background. If this parameter is not supplied, it defaults to true

and the application creation is executed synchronously at the command prompt.

Examples

The following command creates application *app1* in Vserver *vs1* based on the *SAN* application template with *5* LUNs totaling *200 GB* in size and maps the new application LUNs to initiator group *ig1* .

```
cluster1::> application create -vserver vs1 -application app1 -application
-template SAN -parameters os-type:linux,application-
components[0].name:data,application-components[0].total-
size:200GB,application-components[0].lun-count:5,application-
components[0].igroup-name:ig1
```

application delete

Delete an application

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

Description

This command deletes an application.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

The Vserver from which to delete the application.

-application <text> - Application (privilege: advanced)

The name of the application to delete.

[-foreground {true|false}] - Execute in the Foreground (privilege: advanced)

A boolean value that determines if the application deletion executes synchronously at the command prompt in the foreground or asynchronously in the background. If this parameter is not supplied, it defaults to true and the application deletion is executed synchronously at the command prompt.

Examples

The following command deletes application *app1* from Vserver *vs1* .

```
cluster1::> application delete -vserver vs1 -application app1
```

application modify

(DEPRECATED)-Modify properties of an application

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

Description

This command modifies properties of an application. The user must specify the names of one or more existing application components of an existing application and properties to apply to those application components. The command fails and the application remains in its initial state unless all of the application components can be modified at once.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

The Vserver hosting the application.

`-application <text>` - Application (privilege: advanced)

The name of the application to modify.

`-application-components <Application Component>,…` - Application Components (privilege: advanced)

A list of application component names to modify.

`[-storage-service <text>]` - Storage Service (privilege: advanced)

The new storage service to use for the listed application components. The storage-service show command displays the available storage services.

If the current aggregates hosting the application components can accommodate the requested storage service, the application components are modified to use the new storage service.

If the current aggregates hosting the application components cannot accommodate the requested storage service, the command fails and all the specified application components continue with their existing storage services.

Modifying the storage service of application components results in modifying the associated QoS policies to reflect the new storage service. The contents of the application components continue to remain on the same storage elements.

`[-foreground {true|false}]` - Execute in the Foreground (privilege: advanced)

A boolean value that determines if the application modification executes synchronously at the command prompt in the foreground or asynchronously in the background. If this parameter is not supplied, it defaults to true and the application modification is executed synchronously at the command prompt.

Examples

The following command modifies application *`app1`* in Vserver *`vs1`* based on the *`BasicModify`* application template. Application components *`component1`* and *`component2`* are modified to use the *`extreme`* storage service.

```
cluster1::> application modify -vserver vs1 -application app1 -application
-components component1,component2 -storage-service extreme
```

Related Links

- storage-service show

application show-detail

Display application details

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

Description

This command displays storage and protocol elements of applications.

Parameters

{ [-fields <fieldname>,…]

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| [-instance] }

Specifies the display of all available fields for each selected application.

[-vserver <vserver name>] - Vserver

Selects elements of applications belonging to Vservers that match the parameter value.

[-application <text>] - Application

Selects elements of applications with an application name that matches the parameter value.

[-element-type {lun|lun-clone|nas|igroup}] - Type of Object

Selects application elements with type that matches the parameter value. Possible values are *lun* for application provisioned LUNs on SAN applications, *lun-clone* for auxiliary LUNs in SAN applications, *igroup* for protocol elements in SAN applications and *nas* for application provisioned storage and protocol information in NAS applications.

[-application-component <text>] - Application Component

Selects application elements which belong to the application component that matches the parameter value. Only application provisioned storage elements of type *lun* , *nas* can belong to application components.

[-element-name <text>] - Name of Object

Selects application elements with an element name that matches the parameter value.

[-application-template <text>] - Application Template

Selects application elements that belong to applications with templates that match the parameter value.

[-application-template-version <integer>] - Application Template Version

Selects application elements that belong to applications with template version that match the parameter value.

[-application-uuid <UUID>] - Application UUID

Selects application elements that belong to applications whose UUID matches the parameter value.

[-application-component-uuid <UUID>] - Application Component UUID

Selects application elements which belong to the application component whose UUID that matches the parameter value.

[-application-component-storage-service <text>] - Storage Service

Selects application elements which belong to the application component with a storage service that matches the parameter value.

[-nas-element-path <junction path>] - Volume Path

Selects application elements with a NAS element path that matches the parameter value. This field is valid only for elements of type *nas* .

[-nas-element-size {<integer>[KB|MB|GB|TB|PB]}] - Volume Size

Selects application elements with NAS element size that matches the parameter value. This field is valid only for elements of type *nas*

[-cifs-share-name <Share>] - CIFS Share Name

Selects application elements with a CIFS share that matches the parameter value. This field is valid only for elements of type *nas* .

[-cifs-server-fqdn <text>] - CIFS Server Name

Selects application elements with a CIFS Server that matches the parameter value. This field is valid for only elements of type *nas* .

[-cifs-permissions <text>,…] - CIFS User Permissions

Selects application elements with CIFS Permissions that matches the parameter value. This field is valid for only elements of type *nas* .

[-cifs-access-ips <IP Address>,…] - Recommended LIFs for CIFS

Selects application elements with CIFS LIFs that matches the parameter value. This field is valid for only elements of type *nas* . The CIFS access IPs are recommended LIFs for CIFS access. If LIFs local to the volume are available, then these are recommended. If not, two remote LIFs are returned. Data can be accessed over other available LIFs but that is not recommended.

[-nfs-export-policy-name <export policy name>] - Export Policy Name

Selects application elements with an export policy name that matches the parameter value. This field is valid for only elements of type *nas* .

[-nfs-permissions <text>,…] - NFS Permissions

Selects application elements with NFS permissions that matches the parameter value. This field is valid for only elements of type *nas* .

[-nfs-access-ips <IP Address>,…] - Recommended LIFs for NFS

Selects application elements with NFS LIFs that matches the parameter value. This field is valid for only elements of type *nas* . The NFS access IPs are recommended LIFs for NFS access. If LIFs local to the volume are available, then these are recommended. If not, two remote LIFs are returned. Data can be accessed over other available LIFs but that is not recommended.

[-lun <text>] - LUN Name

Selects application elements with a LUN name that matches the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-lun-path <path>] - LUN Path

Selects application elements with LUN path that matches the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-serial <text>] - LUN Serial

Selects application elements with a LUN serial number that matches the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-lun-ids <integer>,…] - LUN IDs

Selects application elements with LUN IDs that match the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-size {<integer>[KB|MB|GB|TB|PB]}] - Element Size

Selects application elements with a LUN size that matches the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-igroups <text>,…] - Igroups

Selects application elements mapped to igroups with names that match the parameter value. This field is valid for only elements of type *lun* and *lun-clone* .

[-initiator <text>,…] - Initiators

Selects application elements mapped to initiators that matches the parameter value. This field is valid for only elements of type *igroup* .

[-is-protection-supported {true|false}] - Is Protection Supported

Selects applications with is-protection-supported flag that matches the parameter value. The is-protection-supported flag indicates if any type of protection is supported for the application.

[-is-protected {true|false}] - Is Application Protected

Selects applications with is-protected flag that matches the parameter value. The is-protected flag indicates if any type of protection is enabled for the application.

[-local-RPO <text>] - Local RPO

Selects applications with local-RPO that matches the parameter value.

[-local-RPO-description <text>] - Local RPO description

Selects applications with local-RPO-descriptions that match the parameter value. The local RPO description provides details about the local RPO associated with the application.

[-remote-RPO <text>] - Remote RPO

Selects applications with remote-RPO that matches the parameter value.

[-remote-RPO-description <text>] - Remote RPO description

Selects applications with remote-RPO-descriptions that match the parameter value. The remote RPO description provides details about the remote RPO associated with the application.

Examples

```
cluster1::> application show-detail -vserver vs0
rupancluster-1::*> application show-detail
==========================================================================
===
                   Vserver: vs0
               Application: my_nas_app
      Application Template: NAS
==========================================================================
===
NAS Storage Element
                    -------------------
                 Component: data
           Storage Service: value
                   Element: my_nas_app_data_1
                      Path: /my_nas_app_data_1
                      Size: 30MB
                CIFS Share: my_nas_app_data_1
               CIFS Server: KTCIFS.
     CIFS User Permissions: Everyone / Read
                 CIFS LIFs: 10.10.10.20
             Export Policy: my_nas_app
           NFS Permissions: 0.0.0.0/0 / ro
                  NFS LIFs: 10.10.10.20

==========================================================================
===
                   Vserver: vs0
               Application: my_san_app
      Application Template: SAN
==========================================================================
===
SAN Storage Element
                    -------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_1
                      Path: /vol/my_san_app_data_group1_1/data_group1_1
             Serial Number: w89ST?JTeR7V
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 0
SAN Storage Element
                    -------------------
                 Component: data_group1
```

```
           Storage Service: value
                   Element: data_group1_2
                      Path: /vol/my_san_app_data_group1_1/data_group1_2
             Serial Number: w89ST?JTeR7W
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 1
SAN Storage Element
                    -------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_3
                      Path: /vol/my_san_app_data_group1_1/data_group1_3
             Serial Number: w89ST?JTeR7X
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 2
SAN Storage Element
                    -------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_4
                      Path: /vol/my_san_app_data_group1_1/data_group1_4
             Serial Number: w89ST?JTeR7Y
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 3
SAN Storage Element
                    -------------------
                 Component: redolog_group1
           Storage Service: value
                   Element: redolog_group1_1
                      Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_1
             Serial Number: w89ST?JTeR7Z
                      Size: 120MB
                   Igroups: igroup2
                   LUN IDs: 0
SAN Storage Element
                    -------------------
                 Component: redolog_group1
           Storage Service: value
                   Element: redolog_group1_2
                      Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_2
             Serial Number: w89ST?JTeR7/
```

```
                    Size: 120MB
                 Igroups: igroup2
                 LUN IDs: 1
Auxiliary LUN
                  -------------------
                 Element: redolog_group1_2_clone
                    Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_2_clone
           Serial Number: w89ST?JTeR7a
                    Size: 120MB
SAN Protocol Info
                  -----------------
                  Igroup: igroup1
              Initiators: iqn.1995-08.com.example3, iqn.1995-
08.com.example4
SAN Protocol Info
                  -----------------
                  Igroup: igroup2
              Initiators: iqn.1995-08.com.example1, iqn.1995-
08.com.example2
10 entries were displayed.
```

The above example displays the details of all applications in Vserver vs0.

```
cluster1: : > application show-detail -vserver vs0 -application my_san_app
Vserver: vs0
             Application: my_san_app
    Application Template: SAN
SAN Storage Element
                  -------------------
               Component: data_group1
         Storage Service: value
                 Element: data_group1_1
                    Path: /vol/my_san_app_data_group1_1/data_group1_1
           Serial Number: w89ST?JTeR7V
                    Size: 60MB
                 Igroups: igroup1, igroup2
                 LUN IDs: 0, 0
SAN Storage Element
                   -------------------
                Component: data_group1
         Storage Service: value
                 Element: data_group1_2
                    Path: /vol/my_san_app_data_group1_1/data_group1_2
           Serial Number: w89ST?JTeR7W
```

```
                    Size: 60MB
                Igroups: igroup1, igroup2
                LUN IDs: 1, 1
SAN Storage Element
                 -------------------
              Component: data_group1
        Storage Service: value
                Element: data_group1_3
                   Path: /vol/my_san_app_data_group1_1/data_group1_3
          Serial Number: w89ST?JTeR7X
                    Size: 60MB
                Igroups: igroup1, igroup2
                LUN IDs: 2, 2
SAN Storage Element
                 -------------------
              Component: data_group1
        Storage Service: value
                Element: data_group1_4
                   Path: /vol/my_san_app_data_group1_1/data_group1_4
          Serial Number: w89ST?JTeR7Y
                    Size: 60MB
                Igroups: igroup1, igroup2
                LUN IDs: 3, 3
SAN Storage Element
                 -------------------
              Component: redolog_group1
        Storage Service: value
                Element: redolog_group1_1
                   Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_1
          Serial Number: w89ST?JTeR7Z
                    Size: 120MB
                Igroups: igroup2
                LUN IDs: 0
SAN Storage Element
                 -------------------
              Component: redolog_group1
        Storage Service: value
                Element: redolog_group1_2
                   Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_2
          Serial Number: w89ST?JTeR7
                    Size: 120MB
                Igroups: igroup2
                LUN IDs: 1
Auxiliary LUN
```

```
                    --------------------
                  Element: redolog_group1_2_clone
                     Path:
/vol/my_san_app_redolog_group1_1/redolog_group1_2_clone
             Serial Number: w89ST?JTeR7a
                     Size: 120MB
SAN Protocol Info
                    -----------------
                    Igroup: igroup1
                Initiators: iqn.1995-08.com.example3, iqn.1995-
08.com.example4
SAN Protocol Info
                    -----------------
                    Igroup: igroup2
                Initiators: iqn.1995-08.com.example1, iqn.1995-
08.com.example2
9 entries were displayed.
```

The above example displays the details of application my_san_app in Vserver vs0.

```
cluster1: : > application show-detail -vserver vs0 -application my_san_app
-application-component data_group1
                   Vserver: vs0
               Application: my_san_app
      Application Template: SAN
SAN Storage Element
                    --------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_1
                      Path: /vol/my_san_app_data_group1_1/data_group1_1
             Serial Number: w89ST?JTeR7V
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 0
SAN Storage Element
                    --------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_2
                      Path: /vol/my_san_app_data_group1_1/data_group1_2
             Serial Number: w89ST?JTeR7W
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 1
```

```
SAN Storage Element
                    -------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_3
                      Path: /vol/my_san_app_data_group1_1/data_group1_3
             Serial Number: w89ST?JTeR7X
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 2
SAN Storage Element
                    -------------------
                 Component: data_group1
           Storage Service: value
                   Element: data_group1_4
                      Path: /vol/my_san_app_data_group1_1/data_group1_4
             Serial Number: w89ST?JTeR7Y
                      Size: 60MB
                   Igroups: igroup1
                   LUN IDs: 3

4 entries were displayed.
```

The above example displays the details elements belonging to application-component data_group1 in application my_san_app in Vserver vs0.

application show-statistics

Display application statistics

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

Description

This command displays dynamic application statistics.

Parameters

{ [-fields <fieldname>,…]

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| [-instance] }

Specifies the display of all available fields for each selected application.

[-vserver <vserver name>] - Vserver

Selects applications belonging to Vservers that match the parameter value.

[-application <text>] - Application

Selects applications with names that match the parameter value.

[-application-component <text>] - Application Component

Selects application components with names that match the parameter value.

[-application-uuid <UUID>] - Application UUID

Selects applications with the unique identifier that match the parameter value.

[-application-component-uuid <UUID>] - Application Component UUID

Selects application components with the unique identifier that match the parameter value.

[-storage-service <text>] - Storage Service

Selects application components with storage services that match the parameter value.

This is the storage service chosen for the application component.

[-space-provisioned {<integer>[KB|MB|GB|TB|PB]}] - Provisioned Space

Selects applications with space provisioned that matches the parameter value.

This is the maximum space that can be consumed. In some cases, the provisioned space can be capped by the space available in the aggregate.

[-space-used {<integer>[KB|MB|GB|TB|PB]}] - Used Space

Selects applications with space used that matches the parameter value.

This is the space that would be actively used by the application or application component.

[-space-used-excluding-reserves {<integer>[KB|MB|GB|TB|PB]}] - Used Space Excluding Reserves

Selects applications with space used excluding reserves that matches the parameter value.

This is the space actively used excluding any space used from reserves.

[-space-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Used Space

Selects applications with logical space used that matches the parameter value.

This is the space actively used if there were no space efficiency technologies being employed.

[-space-reserved-unused {<integer>[KB|MB|GB|TB|PB]}] - Unused Reserved Space

Selects applications with unused reserved space that matches the parameter value.

This is the space that has been reserved but not used yet. It is not available to the user.

[-space-available {<integer>[KB|MB|GB|TB|PB]}] - Available Space

Selects applications with available space that matches the parameter value.

This is the space available to the user for any new writes. It does not include space already reserved for use earlier.

[-space-savings {<integer>[KB|MB|GB|TB|PB]}] - Saved Space

Selects applications with space savings that match the parameter value.

This is the space saved by using space efficiency technologies.

[-iops-per-tb <integer>] - IOPs Per TB

Selects applications with the number of total operations per TB that match the parameter value.

This is the cumulative total of input and output operations divided by the *space-logical-used*.

[-iops <integer>] - Total IOPs

Selects applications with the number of total operations that match the parameter value.

This is the cumulative total of input and output operations for the application.

[-snapshot-reserve {<integer>[KB|MB|GB|TB|PB]}] - Snapshot Reserve

Selects applications with the Snapshot reserve that matches the parameter value.

This is the space reserved exclusively for Snapshot usage. If Snapshot copies use more space than reserved, they will start to use *space-available*.

[-snapshot-used {<integer>[KB|MB|GB|TB|PB]}] - Snapshot Space Used

Selects applications with the space used for Snapshot copies that match the parameter value.

This is the total space used by Snapshot copies.

Examples

```
cluster1::> application show-statistics
                      Application  Storage    Provisioned       Used
Vserver   Application Component    Service    Space            Space
IOPS
--------- ----------- ------------ ---------- ----------- ----------
----------
vs1       app1
                      -            -             213.8GB     2.39GB
353
                      sg1          value*        166.2GB     1.81GB
75
                      sg2          performance   47.50GB     59.2MB
278
          app2
                      -            -                95GB    48.07GB
2872
                      sg3          extreme          95GB    48.07GB
2872
5 entries were displayed.
```

The above example displays all applications in the cluster. The storage service "value" is appended with an asterisk (*) to indicate some volumes are inaccessible to the application component, resulting in inaccurate statistics. The volumes might be offline. When this occurs, the hidden *stats-incomplete* flag is also set.

application show-template

Show the available application templates

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

Description

This command displays information about the application templates that can be used to provision 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.

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

Selects application templates that match the parameter value.

[-purpose {create|modify}] - Purpose (privilege: advanced)

Selects application templates with a purpose that match the parameter value. The purpose of an application template indicates with application commands with which it can be used.

[-protocol {SAN|NAS}] - Access Protocol (privilege: advanced)

Selects application templates with a protocol that match the parameter value. An application provides resources that are accessed through the protocol.

[-description <text>] - Description (privilege: advanced)

Selects application templates with a description that match the parameter value.

[-missing-prerequisites <text>] - Missing Prerequisites (privilege: advanced)

Selects application templates with missing prerequisites that match the parameter value. Missing prerequisites need to be met before the application template can be used to provision an application.

Examples

```
cluster1::*> application show-template
                   Access
Template Name      Protocol Purpose Description
------------------ -------- ------- ---------------------------------
BasicModify        -        modify  Modify an application.
NAS                NAS      create  A generic NAS application.
OracleOnNFS        NAS      create  Oracle using NFS.
OracleOnSAN        SAN      create  Oracle using SAN.
OracleRACOnNFS     NAS      create  Oracle RAC using NFS.
OracleRACOnSAN     SAN      create  Oracle RAC using SAN.
SAN                SAN      create  A generic SAN application.
VSIOnNAS           NAS      create  A VSI application using NAS.
VSIOnSAN           SAN      create  A VSI application using SAN.
9 entries were displayed.
```

application show

Display applications

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

Description

This command displays applications. An application groups storage objects to perform application-granular operations.

Parameters

{ [-fields <fieldname>,…]

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| [-instance] }

Specifies the display of all available fields for each selected application.

[-vserver <vserver name>] - Vserver

Selects applications belonging to Vservers that match the parameter value.

[-application <text>] - Application

Selects applications with names that match the parameter value.

Examples

```
cluster1::> application show
Vserver   Application
--------- -------------------------------------------
vs1       app1
          app2
          app3
          app4
4 entries were displayed.
```

The example above displays all applications in the cluster.

application provisioning commands

application provisioning revert-volume-qos

Convert the adaptive QoS policy to a traditional QoS policy

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

Description

This command prepares application volumes for a revert to the previous version of Data ONTAP. It converts adaptive QoS policies associated with application volumes to traditional QoS policies.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

The Vserver of the volume whose QoS policy is to be converted.

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

The volume whose QoS policy is to be converted.

Examples

The following command reverts the adaptive QoS policy on volume *`app1_group1_1`* in Vserver *`vs1`* to a traditional QoS policy:

```
cluster1::*> application provisioning revert-volume-qos -vserver vs1
-volume app1_group1_1
```

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
```

application snapshot commands

application snapshot create

Create a snapshot for an application or protection group.

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

Description

This command creates a Snapshot copy of a specified protection-group in a given application. Every application instance also acts as a protection-group and can be treated as one.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the application on which the snapshot is to be created.

`-application <text>` - Application

This specifies the name of the application where a Snapshot is to be created.

`[-protection-group <text>]` - Protection Group

This specifies the name of the protection-group where a Snapshot is to be created. This field is optional, and if not specified, the application itself will be treated as the protection-group.

`-snapshot <snapshot name>` - Snapshot

This specifies the name of the Snapshot that is to be created.

`[-snapshot-consistency-type {crash|application}]` - Snapshot Consistency Type

This specifies whether the snapshot being created is application consistent or crash consistent.

`[-comment <text>]` - Comment

This specifies a comment to be added to a Snapshot copy at the time of creation.

Examples

The following example creates a Snapshot named snap1 on a application named myapp1 on a vserver named vs of a protection-group called myapp1.

```
cluster1::> application snapshot create -vserver vs -application myapp1
-protection-group myapp1 -snapshot snap1
```

application snapshot delete

Delete a snapshot belonging to an application or protection group.

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

Description

This command deletes a Snapshot that belongs to a specified protection-group in a given application. Every application instance also acts as a protection-group and can be treated as one.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the application on which the snapshot is to be deleted.

`-application <text>` - Application

This specifies the name of the application from which a Snapshot is to be deleted.

[-protection-group <text>] - Protection Group

This specifies the name of the protection-group where a Snapshot is to be deleted. This field is optional, and if not specified, the application itself will be treated as the protection-group.

-snapshot <snapshot name> - Snapshot

This specifies the name of the Snapshot that is to be deleted.

[-force <true>] - Force Delete

If this parameter is specified, the Snapshot copy is immediately deleted without generating any confirmation messages. Otherwise the operation generates confirmation messages.

Examples

The following example deletes a Snapshot named snap1 on a application named myapp1 on a vserver named vs of a protection-group called myapp1.

```
cluster1::> application snapshot delete -vserver vs -application myapp1
-protection-group myapp1 -snapshot snap1
```

application snapshot restore

Restore an application or protection group to a snapshot.

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

Description

This command restores a Snapshot that belongs to a specified protection-group in a given application. Every application instance also acts as a protection-group and can be treated as one.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver that contains the application on which the Snapshot is to be restored.

-application <text> - Application

This specifies the name of the application that we want to restore from a Snapshot.

[-protection-group <text>] - Protection Group

This specifies the name of the protection-group where a Snapshot is to be restored. This field is optional, and if not specified, the application itself will be treated as the protection-group.

-snapshot <snapshot name> - Snapshot

This specifies the name of the Snapshot that is to be restored.

Examples

The following example restores a Snapshot named snap1 on a application named myapp1 on a vserver named vs of a protection-group called myapp1.

```
cluster1::> application snapshot restore -vserver vs -application myapp1
-protection-group myapp1 -snapshot snap1
```

application snapshot show

Display a list of snapshots in an application or protection group.

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

Description

This command displays snapshots in a given application or protection-group.

Parameters

{ [-fields <fieldname>,…]

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| [-instance] }

Specifies the display of all available fields for each selected application.

[-vserver <vserver name>] - Vserver

Selects snapshots belonging to vservers that match the parameter value.

[-application <text>] - Application

Selects snapshots with application names that match the parameter value.

[-protection-group <text>] - Protection Group

Selects snapshots with protection-groups that match the parameter value.

[-snapshot <snapshot name>] - Snapshot

Selects snapshots with names that match the parameter value.

[-create-time <Date>] - Creation Time

Selects snapshots with creation times that match the parameter value.

[-protection-group-uuid <UUID>] - Protection Group UUID

Selects applications with protection-group unique identifiers that match the parameter value.

[-application-uuid <UUID>] - Application UUID

Selects snapshots with application unique identifiers that match the parameter value.

[-snapshot-consistency-type {crash|application}] - Snapshot Consistency Type

Selects snapshots with snapshot types that match the parameter value. The type field is used to differentiate between application consistent and crash consistent snapshots.

[-comment <text>] - Comment

Selects snapshots with comments that match the parameter value.

[-is-partial {true|false}] - Partial Snapshot

Selects snapshots with partial attributes that match the parameter value. The is-partial field is used to identify if a snapshot does not exist on all application components. A partial snapshot is still a valid snapshot. It does not exist on all application components because application components were added after the snapshot was created.

Examples

```
cluster1::> application snapshot show
Vserver: vs
Application: myapp1
Protection Group: myapp1
Snapshot                    Create Time           Type
------------------------------ ------------------------ -----------
myapp1_snapshot_1               Wed Dec 31 16:16:40 2017 application
```

The example above displays all snapshots in an application protection-group in the cluster.

application volume commands

application volume remove

Remove a volume from an application

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

Description

This command removes a volume from an application. If the volume is the only member volume of an application component, the application component is deleted. If there are no more application components in the application, the application is deleted.

This command is provided to support downgrade and revert of Data ONTAP to versions prior to the support of applications. It should not be used for any other purpose.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

The Vserver of the volume to remove from the application.

-application <text> - Application (privilege: advanced)

The name of the application from which to remove the volume.

-volume <volume name> - Volume (privilege: advanced)

The volume to remove from the application.

Examples

```
cluster1::*> application volume remove -vserver vs1 -application app1
-volume vol1
```

The example above removes volume vol1 in Vserver vs1 from application app1. If this is the only volume in the application, the application is deleted.

application volume show

Display application volumes

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

Description

This command displays the volumes that are part of an application.

Parameters

{ [-fields <fieldname>,…]

Specifies fields that you want included in the output. You can use `-fields ?` to display the available fields.

| [-instance] }

Specifies the display of all available fields for each selected volume.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Selects application volumes belonging to Vservers that match the parameter value.

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

Selects application volumes that belong to an application that matches the parameter value.

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

Selects application volumes whose name matches the parameter value.

[-size {<integer>[KB|MB|GB|TB|PB]}] - Volume Size (privilege: advanced)

Selects application volumes whose size matches the parameter value.

[-state {online|restricted|offline|force-online|force-offline|mixed}] - Volume State (privilege: advanced)

Selects application volumes whose state matches the parameter value.

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

Selects application volumes that belong to application component that matches the parameter value.

Examples

```
cluster1::> application volume show
Vserver   Application  Volume       State
--------- ------------ ------------ -------
vs1       app1
                       vol1         online
                       vol2         online
                       vol3         online
          app2
                       vol4         online
                       vol5         online
                       vol6         online
vs2       app3
                       vol7         online
                       vol8         online
                       vol9         online
          app4
                       vol10        online
                       vol11        online
                       vol12        online
12 entries were displayed.
```

The example above displays the volumes of all applications in the cluster.

```
cluster1::> application volume show -vserver vs1 -application app1
Vserver   Application  Volume       State
--------- ------------ ------------ -------
vs1       app1
                       vol1         online
                       vol2         online
                       vol3         online
3 entries were displayed.
```

The example above displays the volumes of application app1 in Vserver vs1.

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] }

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 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] }

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
```

autobalance volume commands

autobalance volume rebalance show

Display Auto Balance Volume progress for an Infinite Volume

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

Description

The `autobalance volume rebalance show` command displays information about Auto Balance Volume operations for an Infinite Volume. The command output depends on the parameter or parameters specified with the command. The `autobalance volume rebalance show` command is only supported for Infinite Volumes.

Parameters

{ [-fields <fieldname>,…]

This specifies the fields that need to be displayed.

| [-instance] }

If this parameter is specified, the command displays information about all entries.

[-vserver <vserver name>] - Vserver (privilege: advanced)

If this parameter is specified, the command displays information about capacity balancing for each Infinite Volume and storage service on the specified Vserver.

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

If this parameter is specified, the command displays information about capacity balancing for each storage service on the specified Infinite Volume.

[-storage-service <storage service name>] - Storage Service (privilege: advanced)

If this parameter is specified, the command displays information about capacity balancing for the specified storage-service.

[-state <Auto Balance Volume state>] - State (privilege: advanced)

If this parameter is specified, the command displays information about operations in the specified state.

[-progress <text>] - Progress (privilege: advanced)

If this parameter is specified, the command displays information about operations with the specified progress.

[-transferred {<integer>[KB|MB|GB|TB|PB]}] - Amount Transferred (privilege: advanced)

If this parameter is specified, the command displays information about operations with the specified amount already transferred.

[-target {<integer>[KB|MB|GB|TB|PB]}] - Target Amount (privilege: advanced)

If this parameter is specified, the command displays information about operations with the specified target amount of data to transfer.

[-transferred-percent <percent_no_limit>] - Percentage Transferred (privilege: advanced)

If this parameter is specified, the command displays information about operations with the specified percentage of the transfer complete.

Examples

The following example displays information about all operations on the Vserver named vs0:

```
cluster1::*> autobalance volume rebalance show -vserver vs1
Storage                      Percent
Vserver          Volume        Service    State     Target  Transferred
---------------- ------------ ----------- --------- ------- -----------
vs0              repo_vol     -           running   36.44TB          8%
```

The following example displays information about all operations on the *gold* storage service on the Infinite Volume named *repo_vol* on the Vserver named vs1:

```
cluster1::*> autobalance volume show rebalance -vserver vs1 -volume
repo_vol -storage-service gold
Storage                      Percent
Vserver          Volume        Service    State     Target  Transferred
---------------- ------------ ----------- --------- ------- -----------
vs1              repo_managed gold        running   17.22TB         18%
vs1              repo_managed silver      complete  19.25TB        100%
```

autobalance volume rebalance start

Start Auto Balance Volume for an Infinite Volume

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

Description

The `autobalance volume rebalance start` command allows the user to start Auto Balance Volume and rebalance the used data capacity in an Infinite Volume after files are created. Auto Balance Volume moves data between data constituents of an Infinite Volume. If the Infinite Volume uses storage services, Auto Balance Volume moves data between data constituents of a storage service in an Infinite Volume. Auto Balance Volume ensures that all data constituents in an Infinite Volume or all data constituents in a storage service of an Infinite Volume have similar amounts of used data capacity. The `autobalance volume rebalance start` command is only supported for Infinite Volumes.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

This specifies the Vserver on which the Infinite Volume to be rebalanced is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the Infinite Volume to be rebalanced.

-storage-service <storage service name> - Storage Service (privilege: advanced)

If the Infinite Volume uses storage services, the ` storage-service` parameter is required to specify the storage service to be rebalanced. If the Infinite Volume does not use storage services, the ` storage-service` parameter cannot be specified, and the entire Infinite Volume will be rebalanced.

[-timeout <integer>] - Requisition Timeout (seconds) (privilege: advanced)

The maximum number of seconds Auto Balance Volume will permit an operation to continue without moving files, before moving the operation to the *complete* state.

Examples

The following example starts rebalancing used capacity in the gold storage service for an Infinite Volume named vol:

```
cluster1::*> autobalance volume rebalance start
            -vserver vs0 -volume vol -storage-service gold
```

The following example starts rebalancing used capacity in an Infinite Volume named vol:

```
cluster1::*> autobalance volume rebalance start
            -vserver vs1 -volume vol
```

autobalance volume rebalance stop

Stop Auto Balance Volume for an Infinite Volume

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

Description

The `autobalance volume rebalance stop` command allows the user to stop Auto Balance Volume. The `autobalance volume rebalance stop` command is only supported for Infinite Volumes.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

This specifies the Vserver on which the Infinite Volume being rebalanced is located.

-volume <volume name> - Volume Name (privilege: advanced)

This specifies the Infinite Volume being rebalanced.

-storage-service <storage service name> - Storage Service (privilege: advanced)

If the Infinite Volume being rebalanced uses storage services, the ` storage-service` parameter is required to specify the storage service being rebalanced. If the Infinite Volume being rebalanced does not use storage services, the ` storage-service` parameter cannot be specified because the entire Infinite Volume is

being rebalanced.

Examples

The following example stops rebalancing used capacity for the gold storage service for an Infinite Volume named vol:

```
cluster1::*> autobalance volume rebalance stop
             -vserver vs0 -volume vol -storage-service gold
```

The following example stops rebalancing used capacity for an Infinite Volume named vol:

```
cluster1::*> autobalance volume rebalance stop
             -vserver vs1 -volume vol
```

cluster commands

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 system node show-discovered command displays all the nodes discovered on the local network.

Parameters

{ -node-count <integer> - 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> - Cluster IP Address of Node }

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

[-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 Data ONTAP to be added to the cluster. A Data ONTAP best practice is to add nodes to the cluster that are of the same Data ONTAP version as the nodes in the cluster, but that may not always be possible.

Examples

The following example adds a node using `-node-ip` :

```
cluster1::> cluster add-node -node-ip 192.168.1.5
      [Job 22] Job succeeded.
```

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 node. Once you create the cluster, add additional nodes to the cluster by using the cluster join command.

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.

Parameters

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

This parameter is deprecated and may be removed in a future release of Data 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".

`[-node-count <integer>]` - 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.

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
```

Related Links

- cluster join

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 Data 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>` - 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 Data ONTAP to join the cluster. A Data ONTAP best practice is to join nodes to the cluster that are of the same Data ONTAP version as the nodes in the cluster, but that may not always be possible.

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-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 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 Data 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 Data 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

The following example shows the create option of `cluster setup` .

```
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 Data 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 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 Data 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::>
```

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 unjoin

Unjoin or remove a node from the cluster

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

Description

The `cluster unjoin` 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.

`[-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 unjoin. 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 Data ONTAP to be unjoined from the cluster.

Examples

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

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

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 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 Data ONTAP software contain the timezone data that is current at the time of release. If a change is made to the timezone between Data 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 Data 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 {(ftp|http)://(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-by-error state. The update cannot be canceled if it is not 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-by-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.

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}] - 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] }

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] }

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 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 Data 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 Data 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 Data 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 Data 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 updates of MetroCluster configurations and for two-stage upgrades.

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

Creates a report of the steps that occur during the update without actually doing them.

`[-pause-after {none|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 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 two-pack MetroCluster.

Examples

The following example shows the update operation:

```
cluster1::> cluster image update -version 8.3

It can take several minutes to complete validation...
Pre-update Check       Status      Error-Action
---------------------  ---------
-----------------------------------------
CIFS status            OK
Cluster health status  OK
Cluster quorum status  OK
Disk status            OK
High Availability      OK
status
LIF status             OK
LIFs on home node      OK
status
MetroCluster           OK
configuration status
SnapMirror status      OK
Overall Status         OK
10 entries were displayed.
Do you want to continue? {y|n}: y
Starting update...
```

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 Data 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.

Examples

The following example shows the validate operation:

```
cluster1::> cluster image validate -version 8.3

It can take several minutes to complete validation...
Pre-update Check       Status      Error-Action
---------------------  ---------
-----------------------------------------
CIFS status            OK
Cluster health status  OK
Clsuter quorum status  OK
Disk status            OK
High Availability      OK
status
LIF status             OK
LIFs on home node      OK
MetroCluster           OK
configuration status
SnapMirror status      OK
Overall Status         OK
10 entries were displayed.
```

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 a Data ONTAP package file specified by the URL into the cluster. The package is stored in the cluster package respository 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] }

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 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)

[-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.

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.

`[-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.

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.

[-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.

Examples

```
cluster-1::> cluster log-forwarding show
Verify  Syslog
Destination Host         Port   Protocol        Server  Facility
------------------------ ------ --------------- ------  --------
192.168.0.1              514    udp-unencrypted false   user
```

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}] - 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.

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
Guidance for When Encryption Should Be Used: never
```

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
Guidance for When Encryption Should Be Used: never
```

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
               Guidance for when Encryption Should Be Used: never
```

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.

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
Guidance for when Encryption Should Be Used: as-possible
```

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 wish 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 wish 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`

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 Data 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.

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
```

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
```

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
```

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
```

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
```

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
```

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
```

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 that match this parameter value.

[-idle <[<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.

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: interface_reachable true      true            true
                                       node4
             Data: interface_reachable
             ICMP: interface_reachable true      true            true
node2
           cluster2                    node3
             Data: interface_reachable
             ICMP: interface_reachable true      true            true
                                       node4
             Data: interface_reachable
             ICMP: interface_reachable 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: interface_reachable
       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}]` - 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}] - 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.

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.

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
```

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] }

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 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 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 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)

This optionally specifies 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 Data ONTAP select the NTP protocol version.

The default setting is `auto` .

`[-is-preferred {true|false}]` - Is Preferred NTP Server (default: false) (privilege: advanced)

This optionally specifies whether the external NTP server is the primary time source for correcting and adjusting the cluster time. The responses from this source will be 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.

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
```

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 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)

This optionally specifies 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 Data ONTAP select the NTP protocol version.

The default setting is auto.

`[-is-preferred {true|false}]` - Is Preferred NTP Server (default: false) (privilege: advanced)

This optionally specifies whether the external NTP server is the primary time source for correcting and adjusting the cluster time. The responses from this source will be 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.

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
```

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 Data 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 Data 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 the `-instance` 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 Data ONTAP (`true` ') or not on the list (false`').

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
Server               Version
-------------------- -----------
ntp1.example.com     auto
ntp2.example.com     auto
```

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] }

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 specifed 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>] - 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>] - 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>] - 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 Data ONTAP.

Examples

The following example displays the event catalog:

```
cluster1::> event filter show -filter-name filter1
Filter Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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.

`[-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 events. The documentation for `event show-suppression` command describes the suppression algorithms in detail.

> The suppression parameter can disable both autosuppression and duplicate suppression, but timer suppression cannot be disabled.

`[-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:// prefix. HTTPS connections to a proxy are not supported. To specify a URL that contains a question mark, press ESC followed by the "?".

`[-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.

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 route add-destinations
- event destination create
- event config set-proxy-password

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 via 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.

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.

"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.

"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:  -
```

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
```

event destination commands

event destination create

(DEPRECATED)-Create an event destination

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification destination" command set.

The `event destination create` command creates a new event destination. An event destination is a list of addresses that receive event notifications. These addresses can be e-mail addresses, SNMP trap hosts, and syslog servers. Event destinations are used by event routes. Event routes describe which events generate notifications, and event destinations describe where to send those notifications.

When you create a destination, you can add e-mail addresses, SNMP trap hosts, and syslog hosts to the definition of the destination. Once the destination is fully defined, use the event route add-destinations command to associate the destination with event routes so that notifications of those events are sent to the recipients in the destination.

To see the current list of all destinations and their recipients, use the event destination show command.

There are several default destinations provided for your use.

- allevents - A useful destination for all system events, though no events are routed to this destination by default.
- asup - Events routed to this destination trigger AutoSupport(tm). Only use this destination to send notifications to technical support. See `system node autosupport` for more information.
- criticals - A useful destination for critical events though no events are routed to this destination by default.
- pager - A useful destination for all events that are urgent enough to page a system administrator, though no events are routed to this destination by default.
- traphost - The default destination for all SNMP traps. You can also use the system snmp traphost add command to add SNMP recipients to the traphost default destination.

To add recipients to the default destinations, use the event destination modify command.

You should not create a destination that sends events to more than one type of recipient. Use separate destinations for e-mail, SNMP, and syslog activity. Also, use the traphost default destination for all SNMP activity. You must not create any other destination that sends traps to SNMP trap hosts. The traphost default destination is not required to be added to any event route.

Parameters

`-name <text>` - Name

This mandatory parameter specifies name of the event destination to create.

[-mail <mail address>,…] - Mail Destination

Use this parameter to specify one or more e-mail addresses to which event notifications will be sent. For events to properly generate e-mail notifications, the event system must also be configured with an address and mail server from which to send mail. See event config modify for more information.

[-snmp <Remote IP>,…] - SNMP Destination

To send traps to SNMP trap hosts, use this parameter with the host names or IP addresses of those trap hosts.

[-syslog <Remote IP>,…] - Syslog Destination

Use this parameter with the host names or IP addresses of any remote syslog daemons to which syslog entries will be sent.

[-syslog-facility <Syslog Facility>] - Syslog Facility

This parameter optionally specifies a syslog facility with which the syslog is sent. Possible values for this parameter are default, local0, local1, local2, local3, local4, local5, local6, and local7. If you specify the default syslog facility, syslogs are tagged LOG_KERN or LOG_USER.

[-snmp-community <text>] - SNMP Trap Community

To specify an SNMP trap community, use this parameter with that string.

[-hide-parameters {true|false}] - Hide Parameter Values?

Use this parameter with the value "true" to hide event parameters by removing them from event notifications. This is useful to prevent sensitive information from being sent over non-secure channels.

Examples

The following example creates an event destination named support.email that e-mails events to the addresses supportmgr@example.com, techsupport@example.com, and oncall@example.com.

```
cluster1::> event destination create -name support.email -mail
supportmgr@example.com,techsupport@example.com,oncall@example.com
```

This example creates an event destination named support.bucket01 that sends the notifications to a syslog host.

```
cluster1::> event destination create -name support.bucket01 -syslog
loghost.example.com
```

Related Links

- event route add-destinations
- event destination show
- system snmp traphost add
- event destination modify
- event config modify

event destination delete

(DEPRECATED)-Delete an event destination

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification destination" command set.

The `event destination delete` command removes a specified destination from the list of valid destinations. An event destination is a list of addresses that receive event notifications. These addresses can be e-mail addresses, SNMP trap hosts, and syslog servers. Event destinations are used by event routes. Event routes describe which events generate notifications, and event destinations describe where to send those notifications.

Once you delete a destination, you will not be able to add that destination to any event route.

You will not be able to delete a destination if it is in use by any event routes. To remove a destination from all event routes, so that you can delete it, use the event route remove-destinations`-messagename * -destination `*name* command.

There are several default destinations that cannot be deleted:

- allevents - A useful destination for all system events, though no events are routed to this destination by default.
- asup - Events routed to this destination trigger AutoSupport(tm). Only use this destination to send notifications to technical support. See `system node autosupport` for more information.
- criticals - A useful destination for critical events though no events are routed to this destination by default.
- pager - A useful destination for all events that are urgent enough to page a system administrator, though no events are routed to this destination by default.
- traphost - The default destination for all SNMP traps. You can also use the system snmp traphost delete command to delete SNMP recipients from the traphost default destination.

To see the current list of all destinations, use the event destination show command. To add a new destination to the list, use the event destination create command.

Parameters

`-name <text>` - Name

This mandatory parameter specifies the event destination to delete.

Examples

The following example deletes an event destination named manager.pager:

```
cluster1::> event destination delete -name manager.pager
```

Related Links

- event route remove-destinations
- system snmp traphost delete
- event destination show
- event destination create

event destination modify

(DEPRECATED)-Modify an event destination

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification destination" command set.

The `event destination modify` command changes the definition of an existing event destination. An event destination is a list of addresses that receive event notifications. These addresses can be e-mail addresses, SNMP traphosts, and syslog servers. Event destinations are used by event routes. Event routes describe which events generate notifications, and event destinations describe where to send those notifications.

Modifying a parameter writes over the existing value of the parameter. To extend a parameter, make sure to include the current value of that parameter. For instance, to add an e-mail address to a destination, include all of the current e-mail addresses assigned to that destination along with the new address. To see the current definition of a destination, use the event destination show`-name `*name* command.

You must not create a destination that sends events to more than one type of recipient. Use separate destinations for e-mail, SNMP, and syslog activity. Also, use the traphost default destination for all SNMP activity. You should not create any other destination that sends to SNMP traphosts. The traphost default destination is not required to be added to any event route.

Parameters

`-name <text>` - Name

This mandatory parameter specifies name of the event destination to modify.

`[-mail <mail address>,…]` - Mail Destination

Use this parameter to specify one or more e-mail addresses to which event notifications will be sent. For events to properly generate e-mail notifications, the event system must also be configured with an address and mail server from which to send mail. See event config modify for more information.

`[-snmp <Remote IP>,…]` - SNMP Destination

To send traps to SNMP trap hosts, use this parameter with the host names or IP addresses of those trap hosts.

`[-syslog <Remote IP>,…]` - Syslog Destination

Use this parameter with the host names or IP addresses of any remote syslog daemons to which syslog entries will be sent.

[-syslog-facility <Syslog Facility>] - Syslog Facility

This parameter optionally specifies a syslog facility with which the syslog is sent. Possible values for this parameter are default, local0, local1, local2, local3, local4, local5, local6, and local7. If you specify the default syslog facility, syslogs are tagged LOG_KERN or LOG_USER.

[-snmp-community <text>] - SNMP Trap Community

To specify an SNMP trap community, use this parameter with that string.

[-hide-parameters {true|false}] - Hide Parameter Values?

Enter this parameter with the value "true" to hide event parameters by removing them from event notifications. This is useful to prevent sensitive information from being sent over non-secure channels. Enter it with the value "false" to turn off parameter hiding.

Examples

The following example modifies an event destination named snmp.hosts to send events to SNMP trap hosts named traphost1 and traphost2:

```
cluster1::> event destination modify -name snmp.hosts -snmp
traphost1.example.com,traphost2.example.com
```

This example adds the e-mail address of a remote support facility to an existing list of e-mail recipients.

```
cluster1::> event destination show -name support
Name: support
      Mail Destination: support.hq@company.com
      SNMP Destination: -
    Syslog Destination: -
       Syslog Facility: -
   SNMP Trap Community: -
Hide Parameter Values?: -

cluster1::> event destination modify -name support -mail
support.hq@company.com,support.remote@company.com

cluster1::> event destination show -name support
Name: support
      Mail Destination: support.hq@company.com, support.remote@company.com
      SNMP Destination: -
    Syslog Destination: -
       Syslog Facility: -
   SNMP Trap Community: -
Hide Parameter Values?: -
```

Related Links

- event destination show
- event config modify

event destination show

(DEPRECATED)-Display event destinations

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

Description

> This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification destination" command set.

The `event destination show` command displays information about configured event destinations. An event destination is a list of addresses that receive event notifications. These addresses can be e-mail addresses, SNMP trap hosts, and syslog servers. Event destinations are used by event routes. Event routes describe which events generate notifications, and event destinations describe where to send those notifications.

Default destinations:

- allevents - A useful destination for all system events, though no events are routed to this destination by default.
- asup - Events routed to this destination trigger AutoSupport(tm). Only use this destination to send notifications to technical support. See `system node autosupport` for more information.
- criticals - A useful destination for critical events although no events are routed to this destination by default.
- pager - A useful destination for all events that are urgent enough to page a system administrator, though no events are routed to this destination by default.
- traphost - The default destination for all SNMP traps. You can also use the system snmp traphost show command to view SNMP recipients for the traphost default destination.

To add recipients to the default destination, use the event destination modify command.

> While you can use both host names and IP addresses with parameters, only IP addresses are stored. Unless all DNS and reverse-DNS operations complete successfully, IP addresses might appear in command 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.

| [-facility]

Displays only the syslog destinations and syslog facilities.

| [-instance] }

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

[-name <text>] - Name

Selects the destinations that match this parameter value.

[-mail <mail address>,…] - Mail Destination

Selects the destinations that match this parameter value.

[-snmp <Remote IP>,…] - SNMP Destination

Selects the destinations that match this parameter value (SNMP trap hosts).

[-syslog <Remote IP>,…] - Syslog Destination

Selects the destinations that match this parameter value (syslog event notification daemons).

[-syslog-facility <Syslog Facility>] - Syslog Facility

Selects the destinations that match this parameter value. Valid values are: `default` , `local0` , `local1` , `local2` , `local3` , `local4` , `local5` , `local6` , and `local7` .

[-snmp-community <text>] - SNMP Trap Community

Selects the destinations that match this parameter value.

[-hide-parameters {true|false}] - Hide Parameter Values?

Selects the destinations that match this parameter value (`true` selects destinations that do not receive full event parameters, `false` selects destinations that receive full event parameters). Event parameters may be hidden to prevent sensitive information from being sent over non-secure channels.

Examples

The following example displays information about all event destinations:

```
cluster1::> event destination show
Hide
Name              Mail Dest.         SNMP Dest.          Syslog Dest.
Params
---------------- ----------------- ------------------ ------------------
------
allevents         -                 -                  logger.example.com -
asup              -                 -                  -                  -
criticals         oncall            -                  -                  -
                  @example.com
pager             pager@example.com -                  -                  -
support.email     supportmgr        -                  -                  -
                  @example.com,
                  techsupport
                  @example.com,
                  oncall
                  @example.com      -                  -                  -

traphost          -                 th0.example.com,   -                  -
                                    th1.example.com

6 entries were displayed.
```

Related Links

- system snmp traphost show
- event destination modify

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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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

Filter Name Rule     Rule      Message Name          SNMP Trap Type
Severity
            Position Type
----------- -------- --------- --------------------- ---------------
--------
no-info-debug-events
            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.

Examples

The following example creates an event filter named filter1:

```
cluster1::> event filter create -filter-name filter1

cluster1::> event filter show
Filter Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 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 comamnd.

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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
```

```
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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

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.

Examples

The following example displays the event filters:

```
cluster1::> event filter show
Filter Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
default-trap-events
            1        include   *                      *
EMERGENCY, ALERT
            2        include   *                      Standard, Built-in
                                                                  *
            3        exclude   *                      *              *
important-events
            1        include   *                      *
EMERGENCY, ALERT
            2        exclude   *                      *              *
no-info-debug-events
            1        include   *                      *
EMERGENCY, ALERT, ERROR, NOTICE
            2        exclude   *                      *              *
7 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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
default-trap-events
            2        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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
default-trap-events
            1        include   *                      *
EMERGENCY, ALERT
            3        exclude   *                      *               *
important-events
            1        include   *                      *
EMERGENCY, ALERT
            2        exclude   *                      *               *
no-info-debug-events
            1        include   *                      *
EMERGENCY, ALERT, ERROR, NOTICE
            2        exclude   *                      *               *
6 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 Name Rule     Rule      Message Name           SNMP Trap Type
```

```
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 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 seperated 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 *.

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 specifiying 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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
      Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
      Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
      Position Type
---------- -------- --------- ---------------------- ---------------
--------
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.
```

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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
      Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
```

```
      Position Type
---------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
```

```
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
```

[-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 <text>] - 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 mailhistory commands

event mailhistory delete

(DEPRECATED)-Delete an e-mail history record

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

Description

> This command has been deprecated. It may be removed from a future major release of Data ONTAP. Instead, use the "event notification history" command set.

The `event mailhistory delete` command deletes a record from the e-mail history.

To delete a record, you must know which node contains the record, and the record's sequence number. Use the event mailhistory show command to view this information.

Parameters

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

Use this parameter to specify the name of the node that contains the e-mail history record to delete.

`-seqnum <Sequence Number>` - Sequence Number

Use this parameter to specify the sequence number of the e-mail history record to delete.

Examples

The following example deletes all mail-history records on node1:

```
cluster1::> event mailhistory delete -node node1 -seqnum *
```

Related Links

- event mailhistory show

event mailhistory show

(DEPRECATED)-Display a list of e-mail history records

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification history" command set.

The `event mailhistory show` command displays a list of the event notifications that have been e-mailed. The command output depends on the parameters you specify with the command. By default, the command displays basic information about all notification e-mails that were sent.

To display detailed information about a specific mail-history record, run the command with the `-seqnum` 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 mail-history records that match this parameter value.

[-seqnum <Sequence Number>] - Sequence Number

Selects the mail-history records that match this parameter value.

[-message-name <Message Name>] - Message Name

Selects the mail-history records that match this parameter value.

[-address <mail address>,…] - Mail Address

Selects the mail-history records that match this parameter value.

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

Selects the mail-history records that match this parameter value.

[-message <text>] - Alert Message

Selects the mail-history records that match this parameter value (text pattern).

[-previous-time <MM/DD/YYYY HH:MM:SS>] - Previous Transmission Time

Selects the mail-history records that match this parameter value.

[-num-drops-since-previous <integer>] - Number of Drops Since Previous Transmission

Selects the mail-history records that match this parameter value (number of event drops since last transmission).

Examples

The following example displays detailed information about the mail-history record with the sequence number 20520:

```
cluster1::> event mailhistory show -seqnum 20520
    Sequence Number:  20520
       Message Name:  wafl.vol.full
            Address:  admin@example.com
               Time:  10/1/2008 14:06:24
               Node:  node3
      Previous Time:  5/31/2007 00:33:22
 # Drops Since Prev:  0
       Mail Message:  wafl.vol.full: file system on volume
                      vol0@vserver:28558fe3-2462-11da-85ab
                      -000423bacd20 is full
```

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.

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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 destination 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 destination cannot be added or removed.

Examples

The following example shows the modification 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 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.

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 to which event notifications will be sent. 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 syslog server host name or IP address to which syslog entries will be sent.

| -rest-api-url <text> - REST API Server URL

Use this parameter to specify REST API server URL to which event notifications will be 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 an https:// URL is specified, then Data ONTAP verifies the identity of the destination host by validating its certificate. If the Online Certificate Status Protocol (OCSP) is enabled for

EMS, then Data 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.

[-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.

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
```

Related Links

- event config modify
- 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 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 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.

| [-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 an https:// URL is specified, then Data ONTAP verifies the identity of the destination host by validating its certificate. If the Online Certificate Status Protocol (OCSP) is enabled for EMS, then Data 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.

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

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: -
```

Related Links

- security certificate show

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. Note: In the case of a rest-api destination type, OCSP information is not included. It's available in 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 will not be sent to the server.

Examples

The following shows examples of "event notification destination show":

```
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")
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 values: 10.30.40.10 (from "system snmp traphost")
```

Related Links

- 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 Name Rule     Rule      Message Name           SNMP Trap Type
Severity
            Position Type
----------- -------- --------- ---------------------- ---------------
--------
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 route commands

event route add-destinations

(DEPRECATED)-Add destination(s) to an event definition

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

Description

> This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification" command set.

The `event route add-destinations` command adds destinations to an event route. Any existing destinations assigned to the route are not removed.

The destinations you add must already exist. See the documentation for the event destination create command for information about creating destinations. To show all existing destinations and their attributes, use the event destination show command. To remove destinations from an event route, use the event route remove-destinations command.

You can use extended queries with such parameters as `-severity` and `-snmp-support` to specify multiple events that meet certain criteria. See examples below that show how to use extended queries.

Parameters

`-message-name <Message Name>` - Message Name

Specify the message name of the event you are modifying. You can use wildcards to specify a family of events or type of event.

`[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}]` - Severity

Use this optional parameter to specify a set of events that match this parameter value. You must use the `-message-name` parameter with wildcards to specify the family of events or type of events.

`-destinations <Event Destination>,…` - Destinations

Specify a comma-separated list of destinations to which notifications for the named event are sent. These destinations will be added to any existing destinations assigned to this event route.

Examples

The following example specifies that all RAID events go to the destinations named support.email, mgr.email, and sreng.pager:

```
cluster1::> event route add-destinations -message-name raid* -destinations
support.email,mgr.email,sreng.pager
```

The following example specifies that all alert, and emergency events go to the destination named test_dest:

```
cluster1::> event route add-destinations -message-name * -severity <=ALERT
-destinations test_dest
```

The following example specifies that all alert events that support a SNMP trap go to the destination named traphost. In this example, because the `-snmp-support` parameter is specified as part of extended queries, the `-severity` parameter must also be specified in the extended queries:

```
cluster1::> event route add-destinations {-snmp-support true -severity
ALERT} -destinations traphost
```

Related Links

- event destination create
- event destination show
- event route remove-destinations

event route modify

(DEPRECATED)-Modify an event's destination, reporting threshold, or both

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification" command set.

Use the `event route modify` command to modify an event's destination, frequency threshold, and time threshold. The event's destination must already exist; see the documentation for the event destination create command for information about creating destinations. The frequency threshold and time threshold prevent multiple event notifications in a brief period of time.

You can use extended queries with such parameters as `-severity` and `-snmp-support` to specify multiple events that meet certain criteria. See examples provided in the event route add-destinations command manpage that show how to use extended queries.

The frequency threshold specifies the number of times an event occurs before a repeat notification of the event is sent; for instance, a frequency threshold of 5 indicates that a notification is sent every fifth time an event occurs. The time threshold specifies the number of seconds between notifications for an event; for instance, a time threshold of 120 indicates that a notification is sent only if it has been two minutes or more since the last notification for that event was sent.

If both the frequency threshold and time threshold are set, a notification is sent if either threshold is met. For instance, if the frequency threshold is set to 5 and the time threshold is set to 120, and the event occurs more than five times within two minutes, a notification is sent. If both thresholds are set to 0 (zero) or empty ("-" or ""), there is no suppression of multiple event notifications.

Parameters

-message-name <Message Name> - Message Name

Specify the message name of the event you are modifying. You can use wildcards to specify a family of events or type of event.

[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}] - Severity

Use this optional parameter to specify a set of events that match this parameter value. You must use the `-messagename` parameter with wildcards to specify the family of events or type of events.

[-destinations <Event Destination>,…] - Destinations

Use this optional parameter to specify a comma-separated list of destinations to which notifications for the named event are sent. Using this parameter replaces the current list of destinations with the list of destinations you specify. To add or remove individual destinations from the current list, use event route add-destinations or event route remove-destinations .

[-frequencythreshold <integer>] - Number of Drops Between Transmissions

Specifies the number of event notifications that must occur within the `timethreshold` period before a repeat notification is sent.

[-timethreshold <integer>] - Dropping Interval (Seconds) Between Transmissions

If multiple notifications of an event occur within this many seconds, only the first notification is sent. Multiple notifications will be sent during this time period only if the `frequencythreshold` quantity is exceeded.

Examples

The following example modifies all RAID events to send messages to a destination named "support.email", and specify that multiple messages should only be sent if and event occurs more than five times within 60 seconds.

```
cluster1::> event route modify -messagename raid* -destinations
support.email -frequencythreshold 5 -timethreshold 60
```

Related Links

- event destination create
- event route add-destinations
- event route remove-destinations

event route remove-destinations

(DEPRECATED)-Remove destination(s) from an event definition

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification" command set.

The `event route remove-destinations` command can be used to remove existing destinations from an event route. This command removes only the specified destinations from the route, leaving any other destinations assigned to that route.

The named destinations are not deleted, just removed from the specified event route. To delete a destination entirely, use the event destination delete command. To show all existing destinations and their attributes, use the event destination show command.

You can use extended queries with such parameters as `-severity` and `-snmp-support` to specify multiple events that meet certain criteria. See examples provided in the event route add-destinations command manpage that show how to use extended queries.

Parameters

`-message-name <Message Name>` - Message Name

Specify the message name of the event you are modifying. You can use wildcards to specify a family of events or type of event.

`[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}]` - Severity

Use this optional parameter to specify a set of events that match this parameter value. You must use the `-message-name` parameter with wildcards to specify the family of events or type of events.

`-destinations <Event Destination>,…` - Destinations

Specify a comma-separated list of destinations to remove from the event's list of destinations.

Examples

The following example specifies that the destination named "mgr.email" should no longer receive notifications of RAID events.

```
cluster1::> event route remove-destinations -message-name raid*
-destinations mgr.email
```

Related Links

- event destination delete
- event destination show
- event route add-destinations

event route show

(DEPRECATED)-Display event routes

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event catalog" command set.

This command displays information about event routes. Event routes describe which events generate notifications. A route specifies what to watch for, whom to notify, and what to do should a particular event occur. By default, the command displays the following information:

- Message name of the event
- Severity of the event
- Destinations for event notifications
- Frequency threshold for event notifications
- Time threshold for event notifications

To display detailed information about a specific event route, run the command with the `-message-name` parameter, and specify the name of the message. The detailed view adds the following information:

- Full description of the event
- Action to be taken to address the event

You can specify additional parameters to limit output to the information that matches those parameters. For example, to display information only about events with a message name that begins with "raid", run the command with the `-message-name raid*` parameter. You can enter either a specific text string or a wildcard pattern.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 event routes that match this parameter value.

[-severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}] - Severity

Selects the event routes that match this parameter value. Valid values:

- EMERGENCY - Disruption
- ALERT - Single point of failure
- ERROR - Degradation
- NOTICE - Infromation
- INFORMATIONAL - Information
- DEBUG - Debug information

[-action <text>] - Corrective Action

Selects the events that match this parameter value. This parameter is most useful when entered with wildcards. 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. This parameter is most useful when entered with wildcards. The "description" field describes why the event was encountered and what it means.

[-snmp-support {true|false}] - Supports SNMP trap

Selects the event routes that match this parameter value.

[-destinations <Event Destination>,…] - Destinations

Selects the event routes that match this parameter value. A destination is a list of email addresses, SNMP clients, and syslogs.

[-frequencythreshold <integer>] - Number of Drops Between Transmissions

Selects the event routes that match this parameter value (number of events since previous notification).

[-timethreshold <integer>] - Dropping Interval (Seconds) Between Transmissions

Selects the event routes that match this parameter value.

Examples

The following example displays information about all event routes:

```
cluster1::> event route show
                                                  Freq    Time
Message                Severity  Destinations   Threshd Threshd
---------------------- --------- -------------- ------- -------
admin.config.backup.
push.fail              ERROR     allevents,pager      5     120
admin.config.changed   INFO      allevents            0       0
admin.file.deleted     INFO      allevents            0       0
admin.login.failure    INFO      allevents            0       0
admin.software.
commit.failure         ERROR     criticals,allevents  0       0
admin.software.
commit.success         INFO      allevents            0       0
admin.software.
committing             INFO      allevents            0       0
admin.software.
installed              INFO      allevents            0       0
aggrcopy.dst.
autoRestrictMsg        NOTICE    allevents            0       0
aggrcopy.dst.
noMemory               ERROR     pager,admin          4     300
...
```

event snmphistory commands

event snmphistory delete

(DEPRECATED)-Delete an SNMP trap history record

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification history" command set.

The `event snmphistory delete` command deletes an SNMP trap-history record. To delete a record, you will need to know which node generated the event, and you will need to know the sequence number of that event in the trap-history.

Use the event snmphistory show command to display a list of trap-history records and their sequence numbers.

Parameters

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

Use this parameter to specify the name of the node that contains the snmp history record to delete.

`-seqnum <Sequence Number>` - Sequence Number

Use this parameter to specify the sequence number of the SNMP trap-history record to delete.

Examples

The following example deletes all SNMP trap-history records on node1:

```
cluster1::> event snmphistory delete -node node1 -seqnum *
```

Related Links

- event snmphistory show

event snmphistory show

(DEPRECATED)-Display a list of SNMP trap history records

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

Description

This command has been deprecated. It may be removed from a future release of Data ONTAP. Instead, use the "event notification history" command set.

The `event snmphistory show` command displays a list of event notifications that have been sent to SNMP traps. The command output depends on the parameters specified with the command. By default, the command displays general information about all trap-history records.

To display detailed information about a specific trap-history record, run the command with the `-seqnum` 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 trap-history records that match this parameter value (text pattern).

[-seqnum <Sequence Number>] - Sequence Number

Selects the trap-history records that match this parameter value (sequence number).

[-message-name <Message Name>] - Message Name

Selects the trap-history records that match this parameter value.

[-address <text>,…] - SNMP Client Address

Selects the trap-history records that match this parameter value (IP address).

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

Selects the trap-history records that match this parameter value.

[-message <text>] - Alert Message

Selects the trap-history records that match this parameter value (text pattern).

[-previous-time <MM/DD/YYYY HH:MM:SS>] - Previous Transmission Time

Selects the trap-history records that match this parameter value.

[-num-drops-since-previous <integer>] - Number of Drops Since Previous Transmission

Selects the trap-history records that match this parameter value (number of event drops since last transmission).

Examples

The following example displays information about all SNMP trap-history records:

```
cluster1::> event snmphistory show
Seq # Message Name          Address   Node  Time
----- -------------------- --------- ----- ------------------
12481 raid.mirror.restrict  10.0.2.20 node0 4/14/2008 15:11:04
12482 aggrcopy.dst.noMemory 10.0.2.20 node0 4/14/2008 14:52:54
12483 raid.mirror.restrict  10.0.2.21 node1 4/14/2008 14:41:04
```

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.

| **[-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.

| **[-jobuuid] (privilege: advanced)**

Displays the job ID, the job name, the owning Vserver, and the job UUID.

| **[-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] }

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] }

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] }

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 copy 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.

`-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.

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
------------  --------- ---------------------  ----------
--------------------
hourly         Cluster   -                        98644      mirror-hourly
weeklylog      Node      node0                    1501       log-rotation
weeklylog      Node      node1                    1498       log-rotation
weeklylog      Node      node2                    1499       log-rotation
weeklylog      Node      node3                    1500       log-rotation
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.

[-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       Name          Type     Description
------------- ----------- ---------
---------------------------------------
cluster1
              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.

`-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.

-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.

`-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.

[-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       Name        Description
------------- ----------
-------------------------------------------------
cluster1
              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
       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.

-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.

`-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.

`-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.

[-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       Name        Description
------------- ----------
-------------------------------------------------------
cluster1
              rollingdaily
                          Every 8h
```

lun commands

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.

This command is not supported for FlexGroups or Vservers with Infinite Volumes.

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 are */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 {vmware|hyper_v|windows_2008|windows_gpt|windows|linux|xen|solaris|solaris_efi|hpux|aix|netware|openvms} - OS Type

Specifies the OS type for the new LUN. The OS 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 and 2012 systems.
- xen - the LUN stores Xen data

[-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 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 *disabled* .

[-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, Data 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.

[-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.

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. This command is not supported for a Vserver with Infinite Volume.

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 are */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.

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 Snapshots 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` .

This command is not supported for a Vserver with Infinite 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

Specifies the Vserver.

{ [-path <qtree path>] - Volume or Qtree Path

Specifies the path of the root volume or qtree. 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 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 {vmware|hyper_v|windows_2008|windows_gpt|windows|linux|xen|solaris|solaris_efi|hpux|aix|netware|openvms}] - OS Type

Specifies OS type of the LUN. The OS 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 and 2012 systems.
- xen - the LUN stores Xen data

[-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 the 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 -volume vol0 -ostype netware
Virtual                                      Without    With SS
Complete
Server     Volume       Qtree        OS Type  SS Reserve Reserve   SS
Reserve
---------- ------------ ------------ -------- ---------- ---------
----------
vs0        vol0         ""           netware        45MB      45MB
45MB
```

Displays the maximum size of a LUN for the OS type netware.

```
cluster1::> lun maxsize
                                                    Without   With SS
Complete
Vserver    Volume       Qtree        OS Type  SS Reserve   Reserve SS
Reserve
---------- ------------ ------------ -------- ---------- ---------
----------
vs1        vol1         ""           hyper_v     172.6MB   172.6MB
172.6MB
                                     windows_2008
                                                 172.6MB   172.6MB
172.6MB
                                     windows_gpt 172.6MB   172.6MB
172.6MB
                                     windows     172.6MB   172.6MB
172.6MB
                                     linux         178MB     178MB
178MB
                                     xen           178MB     178MB
178MB
                                     solaris       178MB     178MB
178MB
                                     solaris_efi   178MB     178MB
178MB
                                     hpux          178MB     178MB
178MB
                                     aix           178MB     178MB
178MB
                                     netware       178MB     178MB
178MB
                                     openvms       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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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".

[-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.

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.

This command is not supported for a Vserver with Infinite Volume.

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 a LUN. You can resize a LUN that is mapped and online. However, to prevent any potential problems, take the LUN offline before resizing it.

When you reduce the size of the LUN, the data in the LUN could be truncated. You will receive an error message if you reduce the size of the LUN. To avoid this error message, use the force parameter.

When you increase the size of a LUN, the maximum resize size is based on the initial geometry of the LUN and the currently available space in the volume. You will receive an error message if you exceed this limit. The lun show`-instance` command reports the "Maximum Resize Size" for a LUN based on the initial geometry. The lun maxsize command reports the maximum LUN size based on the available space. The maximum size of the LUN is the smaller of the two limits issued by the lun show`-instance` command or the lun maxsize command.

This command is not supported for a Vserver with Infinite Volume.

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 are */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

Overrides any warnings if you are reducing the size of the LUN. If you use this parameter without a value, it is set to true, and the command does not prompt you when reducing the size of a LUN would produce warnings. If you do not use this parameter, the command displays an error if reducing the size of a LUN would create a problem.

[-size <size>] - New Size

Specifies the new size of the LUN.

- c (1 byte)
- w (2 bytes)
- B (512 bytes)
- k (1024 bytes)
- M (k*k bytes)
- G (k*m bytes)
- T (m*m bytes)

Examples

```
cluster1::> lun resize -vserver vs1 -path /vol/vol1/lun1 -size 500M -force
```

Resizes LUN /vol/vol1/lun1 on Vserver vs1 to 500M, overriding all warnings.

```
cluster1::> lun resize -vserver vs1 -path /vol/vol1/lun1 -size +5m

cluster1::> lun show -vserver vs1 -volume vol1
Vserver   Path                            State   Mapped   Type
Size
--------- ------------------------------- ------- -------- --------
--------
vs1
          /vol/vol1/lun1                  online  mapped   linux
15MB
```

Adds 5M of space to LUN /vol/vol1/lun1 for a total of 15MB.

```
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 -5m -f

cluster1::> lun show -vserver vs1 -volume vol1
Vserver   Path                            State   Mapped   Type
Size
--------- ------------------------------- ------- -------- --------
--------
vs1
          /vol/vol1/lun1                  online  mapped   linux
10MB
```

Resizes the LUN /vol/vol1/lun1 from 15MB to 10MB, overriding all warnings.

Related Links

- lun show
- lun maxsize

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 are */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 {vmware|hyper_v|windows_2008|windows_gpt|windows|linux|xen|solaris|solaris_efi|hpux|aix|netware|openvms}] - 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 and 2012 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* select 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* select 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.

[-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 <size>] - Block Size

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.

[-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 Data ONTAP 7-Mode are assigned new serial numbers for use with Clustered Data 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.

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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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 are */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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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 are */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.

> This command is not supported for a Vserver with Infinite 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 (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 are *`/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 are *`/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.

LUNs cannot be copied to or from a Vserver with Infinite Volume.

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 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.

LUNs cannot be copied to or from a Vserver with Infinite Volume.

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 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.

LUNs cannot be copied to or from a Vserver with Infinite Volume.

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 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.

LUNs cannot be copied to or from a Vserver with Infinite Volume.

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 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.

This command is not supported for a Vserver with Infinite 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>] - 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-Admin|Paused-Error|Complete|Destroyed}] - 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-Admin` - the LUN copy job is in Paused By Admin 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.

[-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).

LUNs cannot be copied to or from a Vserver with Infinite Volume.

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 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>[/.snapshot/<snapshot>][/<qtree>]/<lun>.

`[-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, Snapshot copies 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 Snapshot copies 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 copy *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).

This command is not supported for a Vserver with Infinite Volume.

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.

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 port set so the host knows which LIFs or TPGs to access. When you bind a port set to an igroup, the host knows which iSCSI or FCP LIF to access. If you do not bind an igroup to a port set, 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 port set when you use this command. If you attempt to bind a port set to an initiator group that is already bound to an existing port set, the command returns an error. You can only bind an initiator group to one port set at a time.

If the initiator group is bound, use the lun igroup unbind command to unbind the initiator group from the port set. After the initiator group is unbound, you can bind it to another port set.

You can only bind an initiator group to a non-empty port set.

This command is not supported for a Vserver with Infinite Volume.

Parameters

-vserver <Vserver Name> - Vserver Name

Specifies the Vserver.

-igroup <text> - Igroup Name

Specifies the initiator group that you want to bind a port set to.

-portset <text> - Portset Binding Igroup

Specifies the port set 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 port set 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 port set, a host in the igroup can access the LUNs only by connecting to the target ports in the port set.

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 port set to an initiator when you create an initiator group. You can modify the port set 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.

This command is not supported for a Vserver with Infinite Volume.

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 <protocol_enum>]` - 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 igroup.

| [-i, -iscsi <true>] - iSCSI }

Specifies iSCSI as the protocol type of the new igroup.

-t, -ostype {solaris|windows|hpux|aix|linux|netware|vmware|openvms|xen|hyper_v} - 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 port set 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.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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 {solaris|windows|hpux|aix|linux|netware|vmware|openvms|xen|hyper_v}]` - 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 used in a LUN mapping relationship.

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 and associated LUN maps.

This command is not supported for a Vserver with Infinite Volume.

Parameters

-vserver <Vserver Name> - Vserver Name

Specifies the Vserver.

-igroup <text> - Igroup Name

Specifies the initiator group from which you want to remove an initiator.

-initiator <text>,... - Initiators

Specifies the initiator name you want to remove. Use the WWPN, IQN or the alias of the initiator.

[-f, -force <true>] - Force

Forcibly removes an initiator and any associated LUN maps.

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.

This command is not supported for a Vserver with Infinite Volume.

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 <protocol_enum>] - Protocol

Selects the initiator groups that match this parameter value (`FCP` , `iSCSI` , or `mixed`).

[-t, -ostype {solaris|windows|hpux|aix|linux|netware|vmware|openvms|xen|hyper_v}] - 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.

[-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.

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.

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 port set. When you unbind an initiator group from a port set, all of the initiators in the initiator group have access to all target LUNs on all network interfaces.

This command is not supported for a Vserver with Infinite Volume.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-igroup <text>` - Igroup Name

Specifies the initiator group that you want to unbind from the port set.

Examples

```
cluster1::> lun igroup unbind -vserver vs1 -igroup ig1
```

Unbinds the initiator group ig1 from the port set on Vserver vs1.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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 Data 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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite 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 (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               600604803436313364333336538366537  /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: 600604803436313364333336538366537
                          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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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, please see the Data ONTAP SAN Administration Guide.

This command is not supported for a Vserver with Infinite Volume.

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 Data 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 -volume vol1 -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.

Data 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, Data 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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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, please see the Data ONTAP SAN Administration Guide.

This command is not supported for a Vserver with Infinite Volume.

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*

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] }

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 are */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 {solaris|windows|hpux|aix|linux|netware|vmware|openvms|xen|hyper_v}] - Igroup OS Type

Selects the LUN maps for the initiator groups with the OS type that matches the parameter value. The possible OS types are:

- `solaris` - the LUN stores Solaris raw disk in a single-slice partition.
- `windows` - the LUN stores a raw disk type in a single-partition Windows disk using the Master Boot Record (MBR) partitioning style.
- `hpux` - the LUN stores HP-UX data.
- `linux` - the LUN stores a Linux raw disk without a partition table.
- `netware` - the LUN stores NetWare data.
- `vmware` - the LUN stores VMware data.
- `openvms` - the LUN stores Open-VMS data.

- `xen` - the LUN stores Xen data.
- `hyper_v` - the LUN stores Windows Server 2008 or Windows Server 2012 Hyper-V data.

[-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

```
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.
```

The example above lists all of the mappings between LUNs and initiator groups and the LUN ID for each mapping.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite Volume.

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.

This command is not supported for a Vserver with Infinite 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

Selects LUN move operations that match this parameter value.

[-destination-path <path>] - 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-Admin|Paused-Error|Complete|Destroyed}] - 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-Admin` - the LUN move job is in Paused By Admin 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.

[-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.

This command is not supported for a Vserver with Infinite Volume.

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, Snapshot copies 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 Snapshot copies 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.

This command is not supported for a Vserver with Infinite Volume.

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 are */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] }

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 are */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.

[-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.

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 Vsever 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 port set. To create a new port set, use the lun portset create command.

Use the network interface create command to create new LIFs.

This command is not supported for a Vserver with Infinite Volume.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the port set 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 port set.

Examples

```
cluster1::> portset add -vserver vs1 -portset ps1 -port-name lif1
```

Adds port lif1 to port set 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 port set for FCP and iSCSI. The port set name can include a maximum of 96 characters. You can add LIFs to the new port set. If you do not add a LIF to the port set, you create an empty

port set. To add LIFs to an existing port set, use the lun portset add command.

After you create a port set, you must bind the port set to an igroup so the host knows which FC or iSCSI LIFs to access. If you do not bind an igroup to a port set, 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 port set because the initiators in the igroup would have no LIFs to access the LUN.

This command is not supported for a Vserver with Infinite Volume.

Parameters

-vserver <Vserver Name> - Vserver Name

Specifies the Vserver.

-portset <text> - Portset Name

Specifies the name of the new port set. A port set 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 port set.

| [-i, -iscsi <true>] - iSCSI }

Specifies iSCSI as the protocol type of the new port set.

Examples

```
cluster1::> portset create -vserver vs1 -portset ps1 -protocol mixed
```

Creates a port set *ps1* on Vserver *vs1* with the protocol type of *mixed* .

```
cluster1::> portset create -vserver vs1 -portset iscsips -protocol iscsi
```

Creates a port set *iscsips* on Vserver *vs1* with the protocol type of *iscsi* .

```
cluster1::> portset create -vserver vs1 -portset fcppc -protocol fcp
```

Creates a port set *fcppc* on Vserver *vs1* with the protocol type of *fcp* .

```
cluster1::> portset create -vserver vs1 -portset ps2 -protocol mixed -port
-name l11
```

Creates a port set *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 port set. By default, you cannot delete a port set if it is bound to an initiator group. If a port set is bound to an initiator group, you can do one of the following:

- specify the `force` option to unbind the port set from the initiator group and delete the port set.
- use the lun igroup unbind command to unbind the port set from the initiator group. Then you can delete the port set.

This command is not supported for a Vserver with Infinite Volume.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the port set you want to delete.

`[-f, -force <true>]` - Force

Forcibly unbinds the port set from the initiator group.

Examples

```
cluster1::> portset delete -vserver vs1 -portset ps1
```

Deletes port set 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 port set.

You cannot remove the last LIF in a port set if the port set is bound to an initiator group (igroup). To remove the last LIF in a port set, use the lun igroup unbind command to unbind the port set from the igroup. Then you can remove the last LIF in the port set.

This command is not supported for a Vserver with Infinite Volume.

Parameters

`-vserver <Vserver Name>` - Vserver Name

Specifies the Vserver.

`-portset <text>` - Portset Name

Specifies the port set 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 port set.

Examples

```
cluster1::> port set remove -vserver vs1 -portset ps1 -port-name lif1
```

Removes port lif1 from port set 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 port set. By default, the command displays all LIFs in all port sets.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 port sets that match this parameter value.

[-portset <text>] - Portset Name

Selects the port sets that match this parameter value.

[-port-name <port_name>,…] - LIF Or TPG Name

Selects the port sets that match this parameter value.

[-protocol {mixed|fcp|iscsi}] - Protocol

Selects the port sets that match this parameter value.

[-port-count <integer>] - Number Of Ports

Selects the port sets that match this parameter value.

[-igroups <igroup>,…] - Bound To Igroups

Selects the port sets 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 port sets.

```
cluster1::> lun portset show -port-count 0
Vserver   Portset      Protocol Port Names              Igroups
--------- ------------ -------- ----------------------- ------------
vs1       p1           iscsi    -                       -
```

The example above displays the port sets 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 port sets 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 port sets 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] }

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 Data 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.

This command is not supported for a Vserver with Infinite Volume.

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 Data ONTAP 7-Mode system. This allows the admin to recover space from the volume for LUNs that may not be transitioned to clustered Data ONTAP without distrupting LUNs that have transitioned, for example, if the LUN is an unsupported OS type.

This command is not supported for a Vserver with Infinite Volume.

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 Data 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 clustered Data ONTAP

This command is not supported for a Vserver with Infinite 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 (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 <os_enum>] - 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 clustered Data ONTAP. LUNs where `is-transitioned` is *`false`* do not have a serial number assigned for clustered Data ONTAP.

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

Selects the 7-Mode LUNs that match the specified UUID for clustered Data ONTAP. LUNs where `is-transitioned` is *`false`* do not have a UUID assigned for clustered Data 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.

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 for before running the `metrocluster configure-node` *`mynode`* command, except that only one DR group is configured.

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.
```

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 up.

`[-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 configuration of the automatic switchover. Supported values are as follows:

- 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.

The default value is *auso-on-cluster-disaster* . It is set to default value if `metrolcuster configure` command is issued.

The auto-switchover failure domain is set to *auso-disabled* if the `metrocluster unconfigure` command is issued.

This setting only affects the local cluster where the command is run.

| -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 failure on a node when it is disabled because of internal errors.Possible value for this field is true. This option is available in diag mode and should be used only when warning ems message comes. All the nodes in MCC configuration must have this option enabled (the default state) 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
```

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.

[-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 performning 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 metrocluster check run command.

This command displays the high-level verification results for each of the components. If there 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.

Please 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
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
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:
        5 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 synched.

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}] - 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
clusB
                      negotiated-switchover-ready            not-applicable
                      switchback-ready                       not-applicable
                      job-schedules                          ok
                      licenses                               ok
                      periodic-check-enabled                 ok
                      onboard-key-management                 ok
12 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: 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:
12 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 hearbeat 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 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 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 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.

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.

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 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 comprise 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 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 indentifies 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 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.

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.

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.

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 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.

| [-switch]

Displays details of switches connected to all the interconnect adapters.

| [-connectivity-hidden] (privilege: advanced)

Displays additional connectivity information (IP address, Area ID, Port ID) from all the interconnect adapters to the connected nodes.

| [-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 swithover 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 swithover 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 specidifed 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.

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 specificed, 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 specificed, 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 specificed, 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 specificed, 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 specificed, 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 specificed, 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 specificed, only the operation that has the matching job id displayed.

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 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 it’s 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 it’s 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

[-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

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}] - 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.

-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).

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 possibile before falling into its normal mode of behaviour.

[-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 Data 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 Data ONTAP for remote data access between nodes in the same cluster

The default session-type is AsyncMirrorRemote.

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, upto 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 learnt 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] }

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 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 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 Data 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}] - 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}] - 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}] - 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.

[-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|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|auto}] - Maximum Speed

If this parameter is specified, the command displays information only about the FCP target adapters that match the specified maxium 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|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* .

[-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-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>] - SPF 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.

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
                     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
        Received Optical Power: 441.3 (uWatts)
    Is Received Power In Range: true
 SPF Transmitted Optical Power: 600.4 (uWatts)
        Is Xmit Power In Range: true
```

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-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 visiblity 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.

[-unique <integer>] - Unique

A unique index for each zoneset record.

[-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.

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. Logical interfaces using SAN data protocols do not support migration or failover.

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.

`-role {cluster|data|node-mgmt|intercluster|cluster-mgmt}` - Role

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}]` - Data Protocol

Use this parameter to specify the list of data protocols that can be configured on the LIF. The supported protocols are NFS, CIFS, FlexCache, iSCSI, and FCP. NFS, CIFS, and FlexCache are available by default when you create a LIF. If you specify "none", the LIF does not support any data protocols. Also, none, iscsi, or fcp cannot be combined with any other protocols.

The data-protocol field must be specified when the LIF is created and cannot be modified later.

`-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.

`-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.

{ -auto {true|false} - IPv4 Link Local

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.

[-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-parter 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 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.

Logical interfaces for SAN protocols do not support failover. Thus, such interfaces will always show this parameter as `disabled` .

[-firewall-policy <policy>] - Firewall Policy

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.

[-auto-revert {true|false}] - Auto Revert

Use this parameter to specify whether a data LIF is automatically reverted to its home node 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. The default setting is `false` . If you set the value of this parameter to `true` , load balancing migration capability of the data LIF is disabled (the `-allow-lb-migrate` parameter is set to `false`).

Logical interfaces for SAN traffic do not support auto-revet. Thus, this parameter is always `false` on such interfaces.

[-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 Data 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.

Logical interfaces for SAN protocols do not support failover. Thus, this parameter cannot be specified for such interfaces.

[-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.

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 firewall policy is *data* 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 -role data
-home-node node0 -home-port e0c -address 192.0.2.130 -netmask
255.255.255.128 -failover-policy broadcast-domain-wide -firewall-policy
data -auto-revert true
cluster1::> network interface create -vserver vs0 -lif datalif2 -role data
-home-node node0 -home-port e0c -address 3ffe:1::aaaa -netmask-length 64
-failover-policy broadcast-domain-wide -firewall-policy data -auto-revert
true
```

Related Links

- 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 port set, the LIF must be removed from the port set before it can be deleted. To determine if a LIF is in a port set, use the lun portset show command. To remove the LIF from the port set, 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.

Logical interfaces for SAN protocols do not support migration. Attempts to do so will result in an error.

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

Logical interfaces for SAN protocols do not support migration. Attempts to do so will result in an error.

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.

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 port set. To modify these properties, first remove the LIF from the port set. To determine if a LIF is in a port set, use the lun portset show command. To remove the LIF from the port set, 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

`[-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.

> If you change this parameter for a cluster or management LIF, you must reboot the storage system to force the change to take effect.

`[-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.

[-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.

Logical interfaces for SAN protocols do not support failover. Thus, such interfaces always show this parameter as `disabled` .

[-firewall-policy <policy>] - Firewall Policy

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.

[-auto-revert {true|false}] - Auto Revert

Use this parameter to modify whether a data LIF is reverted automatically to its home node 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. The default setting is `false` . If you set the value of this parameter to `true` , the load balancing migration capability of the data LIF is disabled (the `-allow-lb-migrate` parameter is set to `false`).

Logical interfaces for SAN traffic do not support auto-revert. Thus, this parameter is always `false` on such interfaces.

[-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 Data 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.

Logical interfaces for SAN protocols do not support failover. Thus, this parameter cannot be specified for such interfaces.

[-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.

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> - LIF

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.

[-role {cluster|data|node-mgmt|intercluster|cluster-mgmt}] - Role

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}] - Data Protocol

Use this parameter to display information only about logical interfaces that have the enabled data protocols you specify.

[-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.

[-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.

[-subnet-name <subnet name>] - Subnet Name

Use this parameter to display the logical interfaces that matches the subnet name.

[-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>] - Firewall Policy

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 Data 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.

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

Start the cluster check function

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

Description

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 Data 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 Data 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] }

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 Data 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 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] }

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 wish 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.

Interfaces for SAN protocols do not support failover. Such interfaces are not valid failover targets.

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 wish 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 wish 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 remote-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, interfacde group, or VLAN) you specify.

[-broadcast-domain <broadcast domain name>] - 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] }

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 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. Please 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 name>,…] - 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 ] }`

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 ]` }

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] }

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   40:55:39:25:27:c1    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 ]` }

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 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] }

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 Data ONTAP 8.3, or when the node is offline. To change the IPspace of a port once the node is in a Data

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|1000|10000|25000|40000}]` - 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}]` - 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 Data ONTAP 8.3, or when the node is offline. To change the IPspace of a port once the node is in a Data 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.

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. For example, if the network is underperforming due to network problems, the operational speed value can be lower than the administrative setting.

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.

Parameters

{ [-fields <fieldname>,…]

If you specify 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|1000|10000|25000|40000}] - Speed Administrative

Selects the network ports that match this parameter value.

[-speed-oper {auto|10|100|1000|10000|25000|40000}] - Speed Operational

Selects the network ports that match this parameter value.

[-flowcontrol-admin {none|receive|send|full}] - Flow Control Administrative

Selects the network ports that match this parameter value.

[-flowcontrol-oper {none|receive|send|full}] - 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}] - 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 name>] - 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.

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 name> - 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 name>` - 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 name>` - 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 name>` - Layer 2 Broadcast Domain
The merging broadcast domain.

`-into-broadcast-domain <broadcast domain name>` - 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 name>` - Layer 2 Broadcast Domain
The name of the broadcast domain.

`[-mtu <integer>]` - Configured MTU
MTU of the broadcast domain.

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 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 name> - 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 name> - 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 name>] - 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.

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 name>` - Layer 2 Broadcast Domain

The broadcast domain to split.

`-new-broadcast-domain <broadcast domain name>` - 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 {<netport>|<ifgrp>}` - 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.

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 {<netport>|<ifgrp>}` - 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 {<netport>|<ifgrp>}` - 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 {<netport>|<ifgrp>}` - 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 {<netport>|<ifgrp>}]` - 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 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.

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, SnapMirror, 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, SnapMirror, 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
                     Telnet               48  false
                     iSCSI                10  false
11 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] }

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>] - 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 routing-groups commands

network routing-groups create

(DEPRECATED)-Create a routing group

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

Description

> This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups create` command creates a group of static routes. After you have created a routing group, you can add routes to the group by using the network routing-groups route create command.

Parameters

`-vserver <vserver>` - Vserver Name

Specifies the node or Vserver on which the routing group will be created.

`-routing-group <text>` - Routing Group

Specifies the name of the routing group that you want to create.

`-subnet <IP Address/Mask>` - Address/Mask

Specifies the IP address and subnet mask of the routing group's destination. The format for this value is: address, slash ("/"), mask. The example below has `192.0.2.165/24` as a valid value for the -subnet parameter.

`-role {cluster|data|node-mgmt|intercluster|cluster-mgmt}` - Role

Defines the role of the routing group. The routing group can be a cluster, data, node management, intercluster, or cluster management routing group. There is no default.

`[-metric <integer>]` - Metric

Specifies a hop count for the routing group that you are creating. The default is 20.

Examples

The following example creates a routing group for data from the Vserver node1 with an IP address of 192.0.2.165/24 to a destination server with the IP address of 192.0.2.166.

```
cluster1::network routing-groups> create -vserver node1 -routing-group
192.0.2.166 -subnet 192.0.2.165/24 -role data -metric 20
```

Related Links

- network routing-groups route create

network routing-groups delete

(DEPRECATED)-Delete a routing group

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

Description

This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups delete` command deletes a specified group of static routes.

Before you run this command, you must delete any logical interfaces that are using this routing group. Use the network interface delete command to delete any logical interfaces using this group.

Parameters

`-vserver <vserver>` - Vserver Name

Specifies the node or Vserver from which the routing group will be deleted

`-routing-group <text>` - Routing Group

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

Examples

The following example deletes a routing group from the Vserver node1 with an IP address of 192.0.2.165/24.

```
cluster1::network routing-groups> delete -vserver node1 -routing-group
192.0.2.165/24
```

Related Links

- network interface delete

network routing-groups show

(DEPRECATED)-Display routing groups

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

Description

This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups show` command displays a group of static routes. You can view routes originating from specified servers, and routes with specified names, roles, and number of hops.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 the routing groups within the specified vserver.

[-routing-group <text>] - Routing Group

Use this parameter to display the specified routing group.

[-subnet <IP Address/Mask>] - Address/Mask

Use this parameter to display the routing groups within the specified subnet. The format for this value is: address, slash ("/"), mask. The example below has `192.0.2.165/24` as a valid value for the -subnet parameter.

[-role {cluster|data|node-mgmt|intercluster|cluster-mgmt}] - Role

Use this parameter to display the routing groups with the specified role.

[-metric <integer>] - Metric

Use this parameter to display the routing groups with the specified metric.

[-address-family {ipv4|ipv6|ipv6z}] - Address Family

Use this parameter to display the routing groups using the specified IP address family. Only IPv4 and IPv6 non-zoned addresses can be used as value for this parameter. IPv6z addresses should not be used.

Examples

The following example displays a routing group for data from the virtual server node1.

```
cluster1::> network routing-groups show -role data
         Routing
Server   Group     Subnet           Role      Metric
-------- --------- ---------------- --------- -------
node1    d192.0.2.165/24
                   192.0.2.165/24   data      20
node2    d192.0.2.166/24
                   192.0.2.166/24   data      20
2 entries were displayed.
```

network routing-groups route create

(DEPRECATED)-Create a static route

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

Description

This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups route create` command creates a static route within a routing group. You can create routes originating from specified Vservers within a specified routing group, routes with specified gateways, and routes with a specified number of hops.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the node or Vserver on which the route is to be created.

`-routing-group <text>` - Routing Group

Use this parameter to specify the name of the routing group within which you want to create the new route.

`-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. 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 specify the IP address of the gateway server leading to the route's destination.

`[-metric <integer>]` - Metric

Use this parameter to specify the hop count for the route you are creating. The default is 20 hops.

Examples

The following example creates a route within a routing group originating from Vserver node3.

```
cluster1::> network routing-groups route create -vserver node3 -routing
-group d192.0.2.167/24 -destination 0.0.0.0/0 -gateway 10.61.208.1 -metric
10
```

network routing-groups route delete

(DEPRECATED)-Delete a static route

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

Description

This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups route delete` command deletes a static route from a routing group. You can delete routes originating from specified Vservers, and routes within specified routing groups.

Parameters

`-vserver <vserver>` - Vserver Name

Use this parameter to specify the node or Vserver from which the route will be deleted.

`-routing-group <text>` - Routing Group

Use this parameter to specify the name of the routing group within which you want to delete the route.

`-destination <IP Address/Mask>` - Destination/Mask

Use this parameter to specify the IP address and subnet mask of the route you want to delete. 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.

Examples

The following example deletes a route within routing group d192.0.2.167/24 originating from Vserver node3.

```
cluster1::> network routing-groups route delete -vserver node3 -routing
-group d192.0.2.167/24 -destination 0.0.0.0/0
```

network routing-groups route show

(DEPRECATED)-Display static routes

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

Description

> This command has been deprecated and may be removed from a future version of Data ONTAP. Use the "network route" command set to configure routes instead.

The `network routing-groups route show` command displays a group of static routes within one or more routing groups. You can view routes originating from specified servers, routes within specified routing groups, routes with specified gateways, and routes with a specified number of hops.

Parameters

`{ [-fields <fieldname>,…]`

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 the routes within the specified vserver.

[-routing-group <text>] - Routing Group

Use this parameter to display the routes within the specified routing group.

[-destination <IP Address/Mask>] - Destination/Mask

Use this parameter to diplay the routes with the specified destination IP address.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 the routes with the specified gateway.

[-metric <integer>] - Metric

Use this parameter to display the routes with the specified metric.

[-address-family {ipv4|ipv6|ipv6z}] - Address Family

Use this parameter to display the routes using the specified address family. Only IPv4 and IPv6 non-zoned addresses can be used for this parameter. IPv6z addresses should not be used.

Examples

The following example displays information about all routing groups.

```
cluster1::> network routing-groups route show
         Routing
Server   Group     Destination      Gateway          Metric
-------- --------- ---------------- ---------------- -------
node1    d192.0.2.165/24
                   0.0.0.0/0        10.61.208.1      20
node2    d192.0.2.166/24
                   0.0.0.0/0        10.61.208.1      20
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 name>` - 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 name>] - 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 tcpdump commands

network tcpdump show

Show running tcpdump instances

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

Description

The `network tcpdump show` command shows currently running packet traces (via tcpdump) on a matching 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

Use this parameter optionally to show the details of running packet traces on a matching node.

[-port {<netport>|<ifgrp>}] - Port

Use this parameter optionally to show the details of running packet trace on a matching network interface.

Examples

The following example shows the details of running packet traces on nodes "node1" and "node2":

```
cluster1::> network tcpdump show
Node       Port
---------- --------
node1
           e0a
node2
           e0c
```

network tcpdump start

tcpdump start

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

Description

The `network tcpdump start` command starts packet tracing (via tcpdump) with the given parameters.

Parameters

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

Use this parameter to specify the node on which the packet trace should run.

-port {<netport>|<ifgrp>} - Port

Use this parameter to specify the network interface for packet tracing.

[-address <IP Address>] - IP Address

Use this parameter to optionally specify the address for packet tracing.

[-buffer-size <integer>] - Buffer Size in KB

Use this parameter to optionally specify the buffer size for packet tracing. The default buffer size is 4 KB.

[-protocol-port <integer>] - Protocol Port Number

Use this parameter to optionally specify the protocol port number for packet tracing.

[-file-size <integer>] - Trace File Size in MB

Use this parameter to optionally specify the trace file size for packet tracing. The default trace file size is 1 GB.

[-rolling-traces <integer>] - Number of Rolling Trace Files

Use this parameter to optionally specify the number of rolling trace files for packet tracing. The default number of rolling trace files is 2.

Examples

The following example starts packet tracing on node "node1" on Vserver "vs0" with address "10.98.16.164", network interface "e0c", buffer size "10 KB", and protocol port number "10000":

```
cluster1::> network tcpdump start -node node1
    -address 10.98.16.164 -port e0c -buffer-size 10 -protocol-port 10000
```

network tcpdump stop

Stop an active tcpdump trace

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

Description

The `network tcpdump stop` command stops a running packet trace (via tcpdump) on a given network interface.

Parameters

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

Use this parameter to specify the node on which the packet tracing must be stopped.

-port {<netport>|<ifgrp>} - Port

Use this parameter to specify the network interface on which the packet tracing must be stopped.

Examples

The following example stops a packet trace on network interface "e0a" from node "node1":

```
cluster1::> network tcpdump stop -node node1 -port e0a
```

network tcpdump trace delete

Delete a tcpdump tracefile

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

Description

The `network tcpdump trace delete` command deletes the tcpdump trace file from a matching node.

Parameters

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

Use this parameter to delete the tcpdump trace file from a matching node.

`-trace-file <text>` - Trace File

Use this parameter to specify the tcpdump trace file to be deleted.

Examples

The following example deletes the list of tcpdump trace files from node "node1" using wildcard pattern:

```
cluster1::> network tcpdump trace delete -node node1 -trace-file *
```

network tcpdump trace show

Show list of tcpdump trace files

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

Description

The `network tcpdump trace show` command shows the list of tcpdump trace files.

Parameters

`{ [-fields <fieldname>,…]`

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 show the list of traces files of a matching node.

[-trace-file <text>] - Trace File

Use this parameter optionally to show the list of trace files with a matching trace-file name.

Examples

The following example shows the list of trace files on nodes "node1" and "node2":

```
cluster1::> network tcpdump trace show
Node             Trace File
---------------- ------------------------------
node1
                 e0a_20170314_115624.trc0
node2
                 e0c_20170314_115624.trc0
```

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 ONTAP Cloud 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] }

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 ONTAP Cloud 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                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.

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.

Examples

```
cluster1::> network tuning icmp6 show
Accept Redirect Redirect Timeout
Node      ICMPv6          in Seconds
--------- --------------- ----------------
node1
          true            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 segements during slow start.

[-is-rfc3390-enabled {true|false}] - RFC3390 enabled

Enables the rfc3390 feature.

[-is-sack-enabled {true|false}] - SACK supoort 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 tuninng 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 supoort 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] }

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 volumes to the policy. When a volume is assigned to an adaptive policy group, the maximum throughput QoS setting automatically adjusts based on volume used space or volume allocated space. QoS minimum throughput setting is calculated from the expected-iops parameter and volume allocated size. It is set only for the volumes that reside on AFF platforms.

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

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 volume allocated size.

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

Specifies the maximum possible IOPS per TB or GB allocated based on volume allocated size or volume 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.

[-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 volume. 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 volume taking into account storage efficiencies. The default value is *used-space* .

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 volume allocated size. QoS minimum throughput setting is calculated from the expected-iops parameter. It is set only for the volumes 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 volume allocated size or volume 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.

[-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 volume. 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 volume taking into account storage efficiencies. The default value is *used-space* .

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 volume used space or volume 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.

[-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.

[-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
Name         Vserver Wklds  IOPS        IOPS
------------ ------- ------ ----------- ------------
extreme      clus-1  3      6144IOPS/TB 12288IOPS/TB
p1           vs0     0      50IOPS/TB   100IOPS/TB
p2           vs0     1      50IOPS/TB   100IOPS/TB
p3           vs0     0      50IOPS/TB   100IOPS/TB
p4           vs0     0      50IOPS/TB   100IOPS/TB
p5           vs0     0      50IOPS/GB   100IOPS/GB
performance  clus-1  0      2048IOPS/TB 4096IOPS/TB
value        clus-1  0      128IOPS/TB  512IOPS/TB

8 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. QoS policy groups cannot belong to Vservers with Infinite Volume.

`[-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.

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: "100iops" and "100".

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.

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.

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: "100iops" and "100".

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 {preset|user-defined|system-defined|autovolume}] - 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`.

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 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 {preset|user-defined|system-defined|autovolume}] - 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.

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

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

Examples

```
cluster1::> qos statistics characteristics show -iterations 100 -rows 4
Policy Group             IOPS      Throughput Request size  Read
Concurrency
-------------------- -------- --------------- ------------  ----
-----------
-total-                    31      304.00KB/s       10041B    0%
16
_System-Best-Effort        15           0KB/s           0B    0%
0
vol1                       11       44.00KB/s        4096B    0%
40
vol2                        4      256.00KB/s       65536B    0%
14
vs1vol0                     1        4.00KB/s        4096B    0%
4
-total-                    37      808.00KB/s       22361B    2%
3
_System-Best-Effort        15           0KB/s           0B    0%
0
vol2                       12      768.00KB/s       65536B    0%
9
vs1vol0                     8       32.00KB/s        4096B   12%
1
vol1                        2        8.00KB/s        4096B    0%
1
```

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

```
cluster1::> qos statistics characteristics show -iterations 100 -policy
-group pg1
Policy Group             IOPS      Throughput Request size Read
Concurrency
-------------------- -------- --------------- ------------ ----
-----------
-total-                   293        3.02MB/s       10783B  54%
0
pg1                       118      470.67KB/s        4096B 100%
0
-total-                   181      478.14KB/s        2700B  65%
0
pg1                       117      469.33KB/s        4096B 100%
0
-total-                   226      525.78KB/s        2382B  60%
1
pg1                       110      440.00KB/s        4096B 100%
1
-total-                   233        1.67MB/s        7527B  49%
1
pg1                       112      446.67KB/s        4096B 100%
1
```

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

qos statistics latency show

Display latency breakdown data per QoS policy group

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

Description

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

The command displays the following data:

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

- Latency observed per I/O operation in the QoS subsystem (QoS)
- Latency observed per I/O operation for NVRAM transfer (NVRAM)
- Latency observed per I/O operation for Object Store(Cloud) 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.

`[-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       NVRAM      Cloud
------------------- ---------- ---------- ---------- ----------
---------- ---------- ---------- ----------
-total-                110.35ms   110.02ms        0ms   327.00us
0ms         0ms          0ms         0ms
vs1vol0                167.82ms   167.22ms        0ms   603.00us
0ms         0ms          0ms         0ms
vol1                   117.76ms   117.56ms        0ms   191.00us
0ms         0ms          0ms         0ms
vol2                    44.24ms    44.05ms        0ms   190.00us
0ms         0ms          0ms         0ms
-total-                 38.89ms    38.63ms        0ms   256.00us
0ms         0ms          0ms         0ms
vol2                    64.47ms    64.20ms        0ms   266.00us
0ms         0ms          0ms         0ms
vol1                    27.28ms    27.03ms        0ms   253.00us
0ms         0ms          0ms         0ms
vs1vol0                 23.72ms    23.47ms        0ms   249.00us
0ms         0ms          0ms         0ms
-total-                409.81ms   409.65ms        0ms   169.00us
0ms         0ms          0ms         0ms
vol1                   816.92ms   816.80ms        0ms   120.00us
0ms         0ms          0ms         0ms
vol2                   407.88ms   407.66ms        0ms   219.00us
0ms         0ms          0ms         0ms
vs1vol0                  3.68ms     3.49ms        0ms   193.00us
0ms         0ms          0ms         0ms
-total-               1169.00us   107.00us        0ms  1062.00us
0ms         0ms          0ms         0ms
vol2                  1169.00us   107.00us        0ms  1062.00us
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      NVRAM       Cloud
------------------- ---------- ---------- ---------- ----------
---------- ---------- ---------- -----------
-total-                 5.88ms   308.00us        0ms   434.00us
5.14ms        0ms        0ms         0ms
pg1                     5.88ms   308.00us        0ms   434.00us
5.14ms        0ms        0ms         0ms
-total-                 4.17ms   280.00us        0ms   477.00us
3.42ms        0ms        0ms         0ms
pg1                     4.17ms   280.00us        0ms   477.00us
3.42ms        0ms        0ms         0ms
-total-                 4.43ms   274.00us        0ms   656.00us
3.50ms        0ms        0ms         0ms
pg1                     4.43ms   274.00us        0ms   656.00us
3.50ms        0ms        0ms         0ms
-total-                 4.89ms   276.00us        0ms   699.00us
3.92ms        0ms        0ms         0ms
pg1                     4.89ms   276.00us        0ms   699.00us
3.92ms        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.

[-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.

[-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.

`[-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 FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

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

The command displays the following data:

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

- Latency observed per I/O operation in the Network subsystem (Network)
- Latency observed per I/O operation across the internally connected nodes in a Cluster (Cluster)
- Latency observed per I/O operation in the Data management subsystem (Data)
- Latency observed per I/O operation in the Storage subsystem (Disk)
- Latency observed per I/O operation in the QoS subsystem (QoS)
- Latency observed per I/O operation for NVRAM transfer (NVRAM)
- Latency observed per I/O operation for Object Store(Cloud) operations

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

Parameters

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

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

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

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

| -vserver <vserver name> - Vserver Name

Specifies the Vserver to which the volume belongs.

-volume <volume name> - Volume Name }

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

[-iterations <integer>] - Number of Iterations

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

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

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

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

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

Examples

```
cluster1::> qos statistics volume latency show -iterations 100 -rows 3
Workload            ID  Latency    Network  Cluster       Data     Disk
QoS      NVRAM      Cloud
--------------- ------ --------   -------- --------   -------- --------
-------- ---------- ----------
-total-                110.35ms   110.02ms      0ms   327.00us      0ms
0ms         0ms         0ms
vs1vol0             111 167.82ms   167.22ms      0ms   603.00us      0ms
0ms         0ms         0ms
vol1               1234 117.76ms   117.56ms      0ms   191.00us      0ms
0ms         0ms         0ms
vol2                999  44.24ms    44.05ms      0ms   190.00us      0ms
0ms         0ms         0ms
-total-              -   38.89ms    38.63ms      0ms   256.00us      0ms
0ms         0ms         0ms
vol2                999  64.47ms    64.20ms      0ms   266.00us      0ms
0ms         0ms         0ms
vol1               1234  27.28ms    27.03ms      0ms   253.00us      0ms
0ms         0ms         0ms
vs1vol0             111  23.72ms    23.47ms      0ms   249.00us      0ms
0ms         0ms         0ms
-total-              - 409.81ms   409.65ms      0ms   169.00us      0ms
0ms         0ms         0ms
vol1               1234 816.92ms   816.80ms      0ms   120.00us      0ms
0ms         0ms         0ms
vol2                999 407.88ms   407.66ms      0ms   219.00us      0ms
0ms         0ms         0ms
vs1vol0             111   3.68ms     3.49ms      0ms   193.00us      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      NVRAM      Cloud
--------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ----------
-total-              -   455.00us   158.00us        0ms   297.00us
0ms        0ms        0ms        0ms
vs0_vol0-wid1..  15658   428.00us   155.00us        0ms   273.00us
0ms        0ms        0ms        0ms
-total-              -   337.00us   130.00us        0ms   207.00us
0ms        0ms        0ms        0ms
vs0_vol0-wid1..  15658   316.00us   128.00us        0ms   188.00us
0ms        0ms        0ms        0ms
-total-              -   464.00us   132.00us        0ms   332.00us
0ms        0ms        0ms        0ms
vs0_vol0-wid1..  15658   471.00us   130.00us        0ms   341.00us
0ms        0ms        0ms        0ms
-total-              -   321.00us   138.00us        0ms   183.00us
0ms        0ms        0ms        0ms
vs0_vol0-wid1..  15658   302.00us   137.00us        0ms   165.00us
0ms        0ms        0ms        0ms
-total-              -   418.00us   142.00us        0ms   276.00us
0ms        0ms        0ms        0ms
vs0_vol0-wid1..  15658   424.00us   143.00us        0ms   281.00us
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 FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

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

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

```
cluster1::> qos statistics volume performance show -vserver vs0 -volume
vs0_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 FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

```
cluster1::> qos statistics volume resource cpu show -node nodeA
-iterations 100 -rows 3
    Workload           ID   CPU
    -------------- ----- -----
    --total- (100%)     -    9%
    vs0vol1-wid-102   102    5%
    vs0vol2-wid-121   121    2%
    vs2_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 FlexVols and Flexgroup Constituents. Otherwise it will display data for FlexVols and Flexgroups.

Examples

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

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

```
cluster1::> qos statistics volume resource disk show -node local -vserver
vs0 -volume vs0_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.

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

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

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

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

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

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

Examples

```
cluster1::> qos statistics workload characteristics show -iterations 100
-rows 4
Workload            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 Data ONTAP subsystems.

The command displays the following data:

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

- Latency observed per I/O operation in the Storage subsystem (Disk)
- Latency observed per I/O operation in the QoS subsystem (QoS)
- Latency observed per I/O operation for NVRAM transfer (NVRAM)
- Latency observed per I/O operation for Object Store(Cloud) 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.

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

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

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

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

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

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

Examples

```
cluster1::> qos statistics workload latency show -iterations 100 -rows 3
Workload            ID  Latency    Network  Cluster       Data     Disk
QoS      NVRAM      Cloud
--------------- ------ --------   -------- --------   -------- --------
-------- ---------- ----------
-total-                110.35ms   110.02ms      0ms   327.00us      0ms
0ms         0ms         0ms
vs1vol0             111 167.82ms   167.22ms      0ms   603.00us      0ms
0ms         0ms         0ms
vol1               1234 117.76ms   117.56ms      0ms   191.00us      0ms
0ms         0ms         0ms
vol2                999  44.24ms    44.05ms      0ms   190.00us      0ms
0ms         0ms         0ms
-total-              -   38.89ms    38.63ms      0ms   256.00us      0ms
0ms         0ms         0ms
vol2                999  64.47ms    64.20ms      0ms   266.00us      0ms
0ms         0ms         0ms
vol1               1234  27.28ms    27.03ms      0ms   253.00us      0ms
0ms         0ms         0ms
vs1vol0             111  23.72ms    23.47ms      0ms   249.00us      0ms
0ms         0ms         0ms
-total-              - 409.81ms   409.65ms      0ms   169.00us      0ms
0ms         0ms         0ms
vol1               1234 816.92ms   816.80ms      0ms   120.00us      0ms
0ms         0ms         0ms
vol2                999 407.88ms   407.66ms      0ms   219.00us      0ms
0ms         0ms         0ms
vs1vol0             111   3.68ms     3.49ms      0ms   193.00us      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      NVRAM      Cloud
--------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ----------
-total-              -     4.80ms   287.00us        0ms   427.00us
4.08ms        0ms         0ms         0ms
file-test1-wi..   7872     9.60ms   265.00us        0ms   479.00us
8.85ms        0ms         0ms         0ms
file-test1_a-..   6437     8.22ms   262.00us        0ms   424.00us
7.53ms        0ms         0ms         0ms
-total-              -     4.20ms   296.00us        0ms   421.00us
3.48ms        0ms         0ms         0ms
file-test1-wi..   7872     8.70ms   211.00us        0ms   489.00us
8.00ms        0ms         0ms         0ms
file-test1_a-..   6437     6.70ms   297.00us        0ms   464.00us
5.94ms        0ms         0ms         0ms
-total-              -     5.90ms   303.00us        0ms     1.71ms
3.88ms        0ms         0ms         0ms
file-test1-wi..   7872    11.36ms   263.00us        0ms     2.06ms
9.04ms        0ms         0ms         0ms
file-test1_a-..   6437     9.48ms   250.00us        0ms     2.30ms
6.93ms        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      NVRAM      Cloud
   --------------- ------ ---------- ---------- ---------- ----------
---------- ---------- ---------- ----------
   -total-               -   443.00us   273.00us        0ms   170.00us
0ms         0ms         0ms         0ms
   file-test1_b-..    9492   440.00us   272.00us        0ms   168.00us
0ms         0ms         0ms         0ms
   -total-               -   577.00us   313.00us        0ms   264.00us
0ms         0ms         0ms         0ms
   file-test1_b-..    9492   607.00us   316.00us        0ms   291.00us
0ms         0ms         0ms         0ms
   -total-               -   475.00us   291.00us        0ms   184.00us
0ms         0ms         0ms         0ms
   file-test1_b-..    9492   476.00us   293.00us        0ms   183.00us
0ms         0ms         0ms         0ms
   -total-               -   628.00us   284.00us        0ms   344.00us
0ms         0ms         0ms         0ms
   file-test1_b-..    9492   591.00us   281.00us        0ms   310.00us
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.

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

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

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

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

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

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

Examples

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

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

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

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

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

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

qos statistics workload resource cpu show

Display CPU resource utilization data per QoS workload

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

Description

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

The command displays the following data:

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

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

Parameters

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

Selects the QOS workloads that match this parameter value.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Examples

```
cluster1::> qos statistics workload resource cpu show -node nodeA
-iterations 100 -rows 3
    Workload           ID   CPU
    --------------- ----- -----
    --total- (100%)     -    9%
    vs0-wid-102       102    5%
    file-bigvmdk-..   121    2%
    vs2_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.

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

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

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

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

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

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

Examples

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

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

```
cluster1::> qos statistics workload resource disk show -node local
-iterations 100 -rows 2 -policy-group pg1
Workload            ID  Disk Number of HDD Disks
--------------- ------ ----- ------------------
-total-              -    3%                 10
file-test1_a-..   6437    6%                  6
file-test1-wi..   7872    6%                  6
-total-              -    3%                 10
file-test1_a-..   6437    5%                  6
file-test1-wi..   7872    5%                  6
-total-              -    3%                 10
file-test1_a-..   6437    6%                  6
file-test1-wi..   7872    6%                  6
```

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

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

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

qos workload commands

qos workload 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 {preset|user-defined|system-defined|autovolume}] - 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.

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 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 Data 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` .

`[-ontapiget {on|off}]` - Enable auditing of Data ONTAP API get operations

This specifies whether get requests for the Data 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 Data 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
               Get Requests:
               -----------------
    CLI:       off
    ONTAPI:    on
```

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.

| [-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 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.

security certificate commands

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 Data 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 Data ONTAP as an SSL client

`[-subtype <kmip-cert>]` - Certificate Subtype

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

`-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 *SHA1* , *SHA256* , *MD5* , *SHA224* , *SHA384* and *SHA512* when the "FIPS Mode" in "security config" is false. When the "FIPS Mode" is true, the possible values are *SHA224* , *SHA256* , *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 Data 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 Data 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 Data ONTAP as an SSL client

[-subtype <kmip-cert>] - Certificate Subtype

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

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 securely to a certificate authority (CA) via any electronic media, to apply for a digital identity certificate.

Parameters

-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.

[-size <size of requested certificate in bits>] - Size of Requested Certificate in Bits

This specifies the number of bits in the private key. The higher the value, the more secure is the key. The default is 2048. Possible values include *512* , *1024* , *1536* and *2048* .

[-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 *SHA1* , *SHA256* and *MD5* .

Examples

This example creates a certificate-signing request with a 2048-bit private key generated by the SHA256 hashing function 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 generate-csr -common-name www.example.com
-size 2048 -country US -state California -locality Sunnyvale -organization
IT -unit Software
-email-addr web@example.com -hash-function SHA256

Certificate Signing Request :
-----BEGIN CERTIFICATE REQUEST-----
MIIBGjCBxQIBADBgMRQwEgYDVQQDEwtleGFtcGxlLmNvbTELMAkGA1UEBhMCVVMx
CTAHBgNVBAgTADEJMAcGA1UEBxMAMQkwBwYDVQQKEwAxCTAHBgNVBAsTADEPMA0G
CSqGSIb3DQEJARYAMFwwDQYJKoZIhvcNAQEBBQADSwAwSAJBAPXFanNoJApT1nzS
xOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJbmXuj6U3a1woUsb13wfEvQnHVFNci
2ninsJ8CAwEAAaAAMA0GCSqGSIb3DQEBCwUAA0EA6EagLfso5+4g+ejiRKKTUPQO
UqOUEoKuvxhOvPC2w7b//fNSFsFHvXloqEOhYECn/NX9h8mbphCoM5YZ4OfnKw==
-----END CERTIFICATE REQUEST-----
Private Key :
-----BEGIN RSA PRIVATE KEY-----
MIIBOwIBAAJBAPXFanNoJApT1nzSxOcxixqImRRGZCR7tVmTYyqPSuTvfhVtwDJb
mXuj6U3a1woUsb13wfEvQnHVFNci2ninsJ8CAwEAAQJAWt2AO+bW3FKezEuIrQlu
KoMyRYK455wtMk8BrOyJfhYsB20B28eifjJvRWdTOBEav99M7cEzgPv+p5kaZTTM
gQIhAPsp+j1hrUXSRj979LIJJY0sNez397i7ViFXWQScx/ehAiEA+oDbOooWlVvu
xj4aitxVBu6ByVckYU8LbsfeRNsZwD8CIQCbZ1/ENvmlJ/P7N9Exj2NCtEYxd0Q5
cwBZ5NfZeMBpwQIhAPk0KWQSLadGfsKO077itF+h9FGFNHbtuNTrVq4vPW3nAiAA
peMBQgEv28y2r8D4dkYzxcXmjzJluUSZSZ9c/wS6fA==
-----END RSA PRIVATE KEY-----

Note: Please 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 Data ONTAP is a client
- *`client`* - includes a self-signed or CA-signed digital certificate and private key to be used for Data ONTAP as an SSL client

`[-subtype <kmip-cert>]` - Certificate Subtype

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

`[-kmip-server-ip <IP Address>]` - (DEPRECATED)-IPv4 and IPv6 address

This parameter is deprecated and might be removed in the future releases of Data ONTAP.

```
This parameter is applicable only to the ``_kmip-cert_`` subtype. It
specifies the IP address of the KMIP server.
```

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
Please enter Certificate: Press <Enter> when done
-----BEGIN CERTIFICATE-----
MIIB8TCCAZugAwIBAwIBADANBgkqhkiG9w0BAQQFADBfMRMwEQYDVQQDEwpuZXRh
cHAuY29tMQswCQYDVQQGEwJVUzEJMAcGA1UECBMAMQkwBwYDVQQHEwAxCTAHBgNV
BAoTADEJMAcGA1UECxMAMQ8wDQYJKoZIhvcNAQkBFgAwHhcNMTAwNDI2MTk0OTI4
WhcNMTAwNTI2MTk0OTI4WjBfMRMwEQYDVQQDEwpuZXRhcHAuY29tMQswCQYDVQQG
EwJVUzEJMAcGA1UECBMAMQkwBwYDVQQHEwAxCTAHBgNVBAoTADEJMAcGA1UECxMA
MQ8wDQYJKoZIhvcNAQkBFgAwXDANBgkqhkiG9w0BAQEFAANLADBIAkEAyXrK2sry
-----END CERTIFICATE-----
Please 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

Please 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

Please 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

Please 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, Data ONTAP acts as an SSL client.

```
cluster1::> security certificate install -vserver vs0 -type server-ca

Please 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 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>] - Certificate Subtype

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

[-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.

Examples

The examples below display information about digital certificates.

```
cluster1::> security certificate show

Vserver    Serial Number   Common 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
          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 {(ftp|http)://(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 *SHA1* , *SHA256* and *MD5* .

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

Please 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 MD5 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 MD5

Please 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-----
Please 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-----
Please enter a password for pkcs12 file:
Please 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>` - 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>] - 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 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 TLSv1.2 and TLSv1.1). 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 Data ONTAP supports the TLSv1.2, TLSv1.1 and TLSv1 protocols. For backward compatibility, Data ONTAP supports adding SSLv3 to the supported-protocols list when FIPS mode is disabled. Use the `-supported-ciphers` parameter to configure only AES, or AES and 3DES, or disable weak ciphers such as RC4 by specifying *`!RC4`* . By default the supported-cipher setting is *`ALL:!LOW:!aNULL:!EXP:!eNULL`* . This setting means that all supported cipher suites for the protocols are enabled, except the ones with no authentication, no encryption, no exports, and low encryption cipher suites (currently those using 64-bit or 56-bit encryption algorithms). Select a cipher suite which is available with the corresponding selected protocol. An invalid configuration may cause some functionality to fail to operate properly. Refer to "https://www.openssl.org/docs/apps/ciphers.html" published by the OpenSSL software foundation for the correct cipher string syntax. After modifying the security configuration, reboot all the nodes manually.

Parameters

`-interface <SSL>` - FIPS-Compliant Interface (privilege: advanced)

Selects the FIPS-compliant interface. Default 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 {TLSv1.2|TLSv1.1|TLSv1|SSLv3}] - Supported Protocols (privilege: advanced)

Selects the supported protocols for the selected interface. Default is *TLSv1.2,TLSv1.1,TLSv1*

[-supported-ciphers <Cipher String>] - Supported Ciphers (privilege: advanced)

Selects the supported cipher suites for the selected interface. Default is *ALL:!LOW:!aNULL:!EXP:!eNULL* .

Examples

The following command enables FIPS mode in the cluster. (Default setting for FIPS mode is *false*)

```
cluster1::> security config modify -interface SSL -is-fips-enabled true
```

The following command modifies supported protocols to TLSv1.2 and TLSv1.1 in the cluster. (Default setting for supported protocols is *TLSv1.2,TLSv1.1,TLSv1*)

```
cluster1::*> security config modify -interface SSL -supported-protocols
TLSv1.2, TLSv1.1
```

The following command modifies supported ciphers to ALL:!LOW:!aNULL:!EXP:!eNULL:!RC4 in the cluster. (Default setting for supported ciphers is *ALL:!LOW:!aNULL:!EXP:!eNULL*)

```
cluster1::*> security config modify -interface SSL -supported-ciphers
ALL:!LOW:!aNULL:!EXP:!eNULL:!RC4
```

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: TLSv1.2,TLSv1.1,TLSv1
- Supported ciphers: ALL:!LOW:!aNULL:!EXP:!eNULL

The default cipher suites represent all suites for the listed protocols except those that have no authentication, no encryption, no exports, and low encryption (below 64 or 56 bit).

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 ciphers that the cluster will support. When all the nodes in the cluster are updated with the modified settings, the cluster security config ready value will be shown as *yes* .

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 <SSL>] - FIPS-Compliant Interface (privilege: advanced)

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 {TLSv1.2|TLSv1.1|TLSv1|SSLv3}] - Supported Protocols (privilege: advanced)

Displays configurations that match the specified protocols.

[-supported-ciphers <Cipher String>] - Supported Ciphers (privilege: advanced)

Displays the configurations that match the specified supported ciphers.

Examples

The following example shows the default security configurations for a cluster.

```
cluster1::> security config show
          Cluster                                                Cluster
Security
Interface FIPS Mode  Supported Protocols     Supported Ciphers Config
Ready
--------- ---------- ----------------------- -----------------
----------------
SSL       false      TLSv1.2, TLSv1.1, TLSv1 ALL:!LOW:         yes
                                             !aNULL:!EXP:
                                             !eNULL
```

The following example shows the security configuration after FIPS mode has been enabled.

```
cluster1::> security config show
          Cluster                                                 Cluster
Security
Interface FIPS Mode  Supported Protocols     Supported Ciphers Config
Ready
--------- ---------- ----------------------- -----------------
----------------
SSL       true       TLSv1.2, TLSv1.1        ALL:!LOW:         yes
                                              !aNULL:!EXP:
                                              !eNULL:!RC4
```

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 <Application supporting SSL/TLS protocol>,…` - 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
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
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 <Application supporting SSL/TLS protocol>,... - 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
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
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] }

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

[-application <Application supporting SSL/TLS protocol>] - 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)

[-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
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

Display Security Configuration Status

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

Description

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 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] }

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 key-manager commands

security key-manager add

Add a key management server

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

Description

This command adds a key management server at the indicated IP address to its list of four possible active key management servers. The command fails if there are already four key management servers configured. This command is not supported when onboard key management is enabled.

Parameters

`-address <IP Address>` - IP Address

This parameter specifies the IP address of the key management server you want to use to store keys.

`[-server-port <integer>]` - Server TCP Port

This parameter specifies the TCP port on which the key management server will listen for incoming connections.

Examples

The following example adds the key management server with address 10.233.1.98, listening for incoming connections on the default TCP port 5696, to the list of key management servers used by the external key manager:

```
cluster-1::> security key-manager add -address 10.233.1.198
```

The following example adds the key management server with address 10.233.1.98, listening for incoming connections on TCP port 15696, to the list of key management servers used by the external key manager:

```
cluster-1::> security key-manager add -address 10.233.1.198 -server-port
15696
```

security key-manager create-key

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 configured key management servers. The command fails if the configured key management servers are already storing more than 128 AKs. If command fails due to more than 128 keys in cluster, delete unused keys on your key management servers and try the command again. This command is not supported when onboard key management is enabled.

Parameters

`[-key-tag <text>]` - Key Tag

This parameter specifies the key tag that you want 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 key-manager query command 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, 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. Data 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 create-key

Verifying requirements...

Node: node1
Creating authentication key...
Authentication key creation successful.
Key ID: 000000000000000002000000000000100D0F7C2462D626B739FE81B89F29A092F.

Node: node2
Key manager restore operation initialized.
Successfully restored key information.
```

The following example creates an authentication key with key-tag "disk1-key":

```
cluster-1::> security key-manager create-key -key-tag disk1-key

Verifying requirements...

Node: node1
Creating authentication key...
Authentication key creation successful.
Key ID: 000000000000000002000000000000100B8297A6189BC24B9B84C1916ED576857.

Node: node2
Key manager restore operation initialized.
Successfully restored key information.
```

The following example creates an authentication key with a user-specified authentication passphrase:

```
cluster-1::> security key-manager create-key -prompt-for-key true

Enter a new passphrase::

Reenter the passphrase::

Verifying requirements...

Node: node1
Creating authentication key...
Authentication key creation successful.
Key ID: 000000000000000002000000000001006268333F870860128FBE17D393E5083B.

Node: node2
Key manager restore operation initialized.
Successfully restored key information.
```

security key-manager delete-key-database

Deletes the key hierarchy for onboard key manager

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

Description

The `security key-manager delete-key-database` command permanently deletes the onboard key-management configuration from all nodes of the cluster.

Examples

The following example deletes the onboard key-management configuration from all nodes of the cluster:

```
cluster-1::*> security key-manager delete-key-database

Warning: This command will permanently delete all keys from onboard key
management.
Do you want to continue? {y|n}: y
```

security key-manager delete-kmip-config

Deletes the KMIP configuration

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

Description

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 Data 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 successfully deleted from all nodes
of the
cluster. The keys stored by the external KMIP servers cannot be deleted by
Data ONTAP,
and must be deleted by using external tools.
```

security key-manager delete

Delete a key management server

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

Description

This command removes the key management server at the indicated IP address from the list of active key management servers. If the indicated key management server is the sole storage location for any key that is in use by Data ONTAP, you will be unable to remove the key server. This command is not supported when onboard key management is enabled.

Parameters

`-address <IP Address>` - IP Address

This parameter specifies the IP address of the key management server you want to remove from use.

Examples

The following example removes the key server at IP address 10.233.1.198 from the set of configured key management servers:

```
cluster-1::> security key-manager delete -address 10.233.1.198
```

security key-manager prepare-to-downgrade

Disables onboard keymanagement features for unsupported versions

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

Description

The `security key-manager prepare-to-downgrade` command disables the onboard key management features that are not supported in releases prior to ONTAP 9.1.0. The features that are disabled are onboard key management support for Metrocluster configurations, and Volume Encryption (VE).

Examples

The following example disables the onboard key management support for Metrocluster configurations and Volume Encryption (VE):

```
cluster1::*> security key-manager prepare-to-downgrade
```

security key-manager query

Displays the key IDs stored in a key management server.

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

Description

This command displays the IDs of the keys that are stored on the key management servers. This command does not update the key tables on the node. To refresh the key tables on the nodes with the key management server key tables, run the security key-manager restore command. This command is not supported when onboard key management 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.

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

This parameter specifies the name of the node that queries the specified key management servers. If this parameter is not specified, then all nodes will query the specified key management servers.

[-address <IP Address>] - IP Address

This parameter specifies the IP address of the key management server that you want to query.

[-key-id <text>] - Key ID

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.

[-count <integer>] - (DEPRECATED)-Key Server's Total Key Count

The value *count* is deprecated and may be removed in a future release of Data ONTAP. This parameter specifies the total number of keys stored in the key management servers. If you specify this parameter, then the command displays only the key IDs retrieved from the key management servers whose total key count matches the specified count number.

[-restored {yes|no}] - Key/Key ID Pair Present in Node's Key Table?

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 'yes' 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 'no' 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-manager-server-status {available|not-responding|unknown}] - Command Error Code

This parameter specifies the connectivity status of the key management server. If you specify this parameter, then the command displays only the key IDs retrieved from the key management servers with specified status.

Examples

The following example shows all the keys on all configured key servers, and whether those keys have been restored for all nodes in the cluster:

```
cluster-1::> security key-manager query
Node: node1
   Key Manager: 10.0.0.10
 Server Status: available

Key Tag                                  Key Type  Restored
---------------------------------------  --------  --------
node1                                    NSE-AK    yes
    Key ID:
000000000000000002000000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
301a4e57-9efb-11e7-b2bc-0050569c227f  VEK       yes
    Key ID:
000000000000000002000000000005004d03aca5b72cd20b2f83eae1531c605e0000000000
000000
Node: node2
   Key Manager: 10.0.0.10
 Server Status: available

Key Tag                                  Key Type  Restored
---------------------------------------  --------  --------
node1                                    NSE-AK    yes
    Key ID:
000000000000000002000000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
301a4e57-9efb-11e7-b2bc-0050569c227f  VEK       no
    Key ID:
000000000000000002000000000005004d03aca5b72cd20b2f83eae1531c605e0000000000
000000

If any listed keys have "no" in the "Restored" column, run "security key-
manager
restore" to restore those keys.
```

The following example shows all keys stored on the key server with address "10.0.0.10" from node "node1" with key-tag "node1":

```
cluster-1::> security key-manager query -address 10.0.0.10 -node node1
-key-tag node1
Node: node1
   Key Manager: 10.0.0.10
 Server Status: available

Key Tag                                  Key Type  Restored
---------------------------------------  --------  --------
node1                                    NSE-AK    yes
    Key ID:
000000000000000020000000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000

If any listed keys have "no" in the "Restored" column, run "security key-
manager
restore" to restore those keys.
```

The following example shows the Volume Encryption Key (VEK) with key-tag (i.e., volume UUID) "301a4e57-9efb-11e7-b2bc-0050569c227f" on nodes where that key has not been restored:

```
cluster-1::*> security key-manager query -key-type VEK -key-tag 301a4e57-
9efb-11e7-b2bc-0050569c227f -restored no
Node: node2
   Key Manager: 10.0.0.10
 Server Status: available

Key Tag                                  Key Type  Restored
---------------------------------------  --------  --------
301a4e57-9efb-11e7-b2bc-0050569c227f     VEK       no
    Key ID:
000000000000000020000000000005004d03aca5b72cd20b2f83eae1531c605e0000000000
000000

If any listed keys have "no" in the "Restored" column, run "security key-
manager restore" to restore those keys.
```

Related Links

- security key-manager restore

security key-manager restore

Restore the authentication key and key ID pairs from the key management servers.

Availability: This command is available to *cluster* 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. This command is not supported when onboard key management 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.

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

This parameter specifies the name of the node that is to load the key IDs into its internal key table. If not specified, all nodes retrieve keys into their internal key table.

[-address <IP Address>] - IP Address

If this parameter is specified, the command restores only from key management server at the specified IP address. If not specified the command restores from all available key management servers.

[-key-tag <text>] - Key Tag

This parameter specifies the value associated with the key ID pair at the time of their creation. If specified, restore only key ID pairs associated with the specified key tag. If not specified, all key ID pairs for the cluster are retrieved.

[-key-ids <text>,…] - Authentication Key ID

If this parameter is specified, the command restores only the specified key IDs.

[-count <integer>] - AK/Key ID Pair Count

The value *count* is deprecated and may be removed in a future release of Data ONTAP. This parameter specifies the total number of keys stored in the key management servers. If this parameter is specified, then the command displays only the key IDs retrieved from the key management servers whose total key count matches the specified count number.

[-key-manager-server-status {available|not-responding|unknown}] - Command Error Code

This parameter specifies the connectivity status of the key management server. If you specify this parameter the command displays only the key IDs retrieved from key management servers with specified status.

Examples

The following command restores keys that are currently on a key server but are not stored within the key tables on the cluster:

```
cluster-1::> security key-manager restore
Node: node1
   Key Manager: 10.0.0.10
 Server Status: available

Key IDs
 ----------------------------------------------------------
00000000000000000200000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
00000000000000000200000000005004d03aca5b72cd20b2f83eae1531c605e0000000000
000000
Node: node2
   Key Manager: 10.0.0.10
 Server Status: available

Key IDs
 ----------------------------------------------------------
00000000000000000200000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
00000000000000000200000000005004d03aca5b72cd20b2f83eae1531c605e0000000000
000000
```

The following loads any keys that exist on the key servers with IP address 10.0.0.10 with key-tag "node1" that are not currently stored in key tables of the nodes in the cluster. In this example, a key with that key-tag was missing from two nodes in the cluster:

```
cluster-1::> security key-manager restore -address 10.0.0.10 -key-tag
node1
Node: node1
   Key Manager: 10.0.0.10
 Server Status: available

Key IDs
 ----------------------------------------------------------
00000000000000000200000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
Node: node2
   Key Manager: 10.0.0.10
 Server Status: available

Key IDs
 ----------------------------------------------------------
00000000000000000200000000001001d71f3b2468d7e16a6e6972d3e6645200000000000
000000
```

security key-manager setup

Configure key manager connectivity

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

Description

The `security key-manager setup` command enables you to configure key management. Data 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 onboard key management, 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 onboard key management configured, you will run this command to refresh the missing keys.

For onboard key management 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 onboard key management when a refresh operation is required (see command description). This parameter is ignored when configuring external key management and during the initial setup of onboard key management.

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 onboard key management? {yes, no} [yes]: no
Would you like to configure the KMIP server environment? {yes, no} [yes]:
yes
```

The following example creates a configuration for onboard 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 onboard key management? {yes, no} [yes]: yes
Enter the cluster-wide passphrase for onboard key management. To continue
the
configuration, enter the passphrase, otherwise type "exit":
Re-enter the cluster-wide passphrase:
After configuring onboard key management, 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 update-passphrase

security key-manager show

Display key management servers

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

Description

This command displays the key management servers configured on the cluster. This command is not supported when onboard key management 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.

| [-status]

If you specify this parameter, the command displays the status of each key management server.

| [-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 you want to retrieve key management server status for. If parameter is not specified, all nodes will retrieve the key management servers status.

[-address <IP Address>] - IP Address

Shows only a key management server registered with the input address. It is also possible to show multiple key management servers.

[-server-port <integer>] - Server TCP Port

If you specify this parameter, the command displays only key servers listening on this port.

Examples

The following example lists all configured key management servers:

```
cluster-1::> security key-manager show

Node                    Registered Key Manager
----------------------  --------------------------
node1                   10.225.89.33
node2                   10.225.89.33
```

The following example lists all configured key management servers, the TCP port on which those servers are expected to listen for incoming KMIP connections, and their server status:

```
cluster-1::> security key-manager show -status

Node                    Port    Registered Key Manager       Status
----------------------  ------  --------------------------
---------------
node1                   5696    10.225.89.33                 available
node2                   5696    10.225.89.33                 available
```

security key-manager update-passphrase

Update cluster-wide passphrase

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

Description

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 onboard key management. 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.

Examples

The following example updates the cluster-wide passphrase used for onboard key management:

```
cluster-1::*> security key-manager 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 backup show"
command.
```

Related Links

- security key-manager setup

security key-manager backup show

Show salt and wrapped keys as a hex dump

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

Description

This command displays the backup information for onboard key management, which would be used to recover the cluster in case of catastrophic situations. The information displayed is for the cluster as a whole (not individual nodes). This command is not supported for an external key management configuration.

Examples

The following example displays the onboard key management backup data for the cluster:

```
cluster-1::> security key-manager backup show
-------------------------BEGIN BACKUP-------------------------
TmV0QXBwIEtleSBCbG9iAAEAAAAEAAAAcAEAAAAAADuD+byAAAAACEAAAAAAAAA
QAAAAAAAAABvOlH0AAAAAMh7qDLRyH1DBz12piVdy9ATSFMT0C0TlYFss4PDjTaV
dzRYkLd1PhQLxAWJwOIyqSr8qY1SEBgm1IWgE5DLRqkiAAAAAAAAACgAAAAAAAAA
3WTh7gAAAAAAAAAAAAAAAAIAAAAAAAgAZJEIWvdeHr5RCAvHGclo+wAAAAAAAAAA
IgAAAAAAAAoAAAAAAAAAEOTcR0AAAAAAAAAAAAAAACAAAAAAAJAGr3tJA/
LRzUQRHwv+1aWvAAAAAAAAAACQAAAAAAAAgAAAAAAAACdhTcvAAAAAJ1PXeBf
ml4NBsSyV1B4jc4A7cvWEFY6lLG6hc6tbKLAHZuvfQ4rIbYAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAA
-------------------------END BACKUP-------------------------
```

security key-manager certificate update

(DEPRECATED)-Update key manager SSL certificates

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 of Data ONTAP.

This command updates the SSL/TLS certificate in-place without requiring the original SSL/TLS certificate to be

deleted. This command is not supported when onboard key management is enabled.

Parameters

`-type {client|server}` - (DEPRECATED)-SSL Certificate Type

This parameter is either "client" or "server". If "client", the internal client certificate is replaced. If "server", the internal server certificate is replaced.

`[-address <IP Address>]` - (DEPRECATED)-Key Manager IP Address

This parameter updates the key manager server certificate for a particular key management server at the given IP address.

Examples

The following example is for updating a server certificate:

```
cluster-1::> security key-manager certificate update -type server -address
10.232.186.8

Node: node1
Key manager 10.232.186.8  certificate-authority certificate will be
updated.
Update successful.
Node: node2
Key manager 10.232.186.8  certificate-authority certificate will be
updated.
Update successful.
```

The following example is for updating a client certificate:

```
cluster-1::> security key-manager certificate update -type client

Node: node1
The system client certificate registered with key manager will be updated.
Update successful.
Node: node2
The system client certificate registered with key manager will be updated.
Update successful.
```

security key-manager external boot-interfaces modify

Modify external key manager logical interfaces

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

Description

This command enables cluster administrators to modify the IP address and route information that the external key manager uses at boot time to restore keys from external key servers.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to modify information on the node that you specify.

`-address-type {ipv4|ipv6|ipv6z}` - Address Type (privilege: advanced)

Use this parameter to modify information for the address-type that you specify.

`[-address <IP Address>]` - Local Interface Address (privilege: advanced)

Use this parameter to modify the IP address that the system will use at boot time to restore keys from external key servers. This parameter implies `-override-default true` .

{ `[-netmask <IP Address>]` - Network Mask (privilege: advanced)

Use this parameter to modify the IP netmask that the system will use at boot time to restore keys from external key servers. This parameter can be used only with address-type ipv4. This parameter implies `-override-default true` .

| `[-netmask-length <integer>]` - Bits in Network Mask (privilege: advanced) }

Use this parameter to modify the IP netmask length that the system will use at boot time to restore keys from external key servers. This parameter implies `-override-default true` .

`[-gateway <IP Address>]` - Gateway (privilege: advanced)

Use this parameter to modify the IP gateway that the system will use at boot time to restore keys from external key servers. This parameter implies `-override-default true` .

`[-port {<netport>|<ifgrp>}]` - Network Port (privilege: advanced)

Use this parameter to modify the port that the system will use at boot time to restore keys from external key servers. The value that you specify cannot be a vlan or ifgrp port. This parameter implies `-override -default true` .

`[-override-default {true|false}]` - Override Default Setting? (privilege: advanced)

Use this parameter to modify the system's selection of boot time IP address and route information. When this value is `false` , the system will use the information associated with a node management LIF. When this value is `true` , then the administrator has chosen to override the defaults.

Examples

The following shows how to modify the port used by node "node2" at boot time to restore keys from external IPv4 key servers. In the example, IPv6 is not enabled in the cluster, so the `-address-type` parameter defaults to ipv4.

```
cluster-1::*> security key-manager external boot-interfaces modify -node
node2 -port e0d
```

The following example shows how to modify the IP address and gateway parameters used by node "node1" at boot time to restore keys from external IPv6 key servers.

```
cluster-1::*> security key-manager external boot-interfaces modify -node
node1 -address-type ipv6 -address fd20:8b1e:b255:814e:749e:11a3:3bff:5820
-gateway fd20:8b1e:b255:814e::1
```

security key-manager external boot-interfaces show

Show external key manager logical interfaces

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

Description

This command enables cluster administrators to view the IP address and route information that the external key manager uses at boot time to restore keys from external 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 (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the node that you specify.

[-address-type {ipv4|ipv6|ipv6z}] - Address Type (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the address-type that you specify.

[-address <IP Address>] - Local Interface Address (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the IP address that you specify.

[-netmask <IP Address>] - Network Mask (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the network mask that you specify.

[-netmask-length <integer>] - Bits in Network Mask (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the network mask length that you specify.

[-gateway <IP Address>] - Gateway (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the gateway that you specify.

[-port {<netport>|<ifgrp>}] - Network Port (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information for the port that you specify.

[-override-default {true|false}] - Override Default Setting? (privilege: advanced)

Use this parameter to display information only about boot-time IP address and route information with the override-default setting that you specify.

Examples

The following example shows how to display the IP address and route information that the external key manager uses at boot time to restore keys. In the example, IPv6 is not enabled in the cluster and, as a result, the command displays information for only the IPv4 address-type. he override-default value is false for all rows, which indicates that the system automatically configured the values based on the node management LIF configuration on the nodes.

```
cluster-1::*> security key-manager external boot-interfaces show
     Address Network                                 Override
Node      Type    Address/Mask       Gateway         Port  Default?
--------- ------- ------------------ --------------- ----- --------
node1
     ipv4    10.224.113.159/24  10.224.113.1    e0M   false
node2
     ipv4    10.224.113.160/24  10.224.113.1    e0M   false
2 entries were displayed.
```

The following example shows how to display the IP address and route information that the external key manager uses at boot time to restore keys. In the example, IPv6 is enabled in the cluster and, as a result, the command displays information for both the IPv4 and IPv6 address-types. The override-default value is false for most rows, which indicates that the system automatically configured the values based on the node management LIF configuration on the nodes. The override-default value for node1 and address-type ipv4 is true, which indicates an administrator has used the security key-manager external boot-interfaces modify command to override one or more fields, and that the values may differ from the corresponding node management LIF.

```
cluster-1::*> security key-manager external boot-interfaces show
         Address Network                                    Override
Node     Type    Address/Mask       Gateway         Port   Default?
-------- ------- ------------------ --------------- ------ --------
node1
     ipv4    10.224.113.159/24  10.224.113.1    e0d    true
         ipv6    fd20:8b1e:b255:814e:32bd:f35c:832c:5a09/64
                                    fd20:8b1e:b255:814e::1
                                                    e0M    false
node2
     ipv4    10.224.113.160/24  10.224.113.1    e0M    false
         ipv6    fd20:8b1e:b255:814e:749e:11a3:3bff:5820/64
                                    fd20:8b1e:b255:814e::1
                                                    e0M    false
4 entries were displayed.
```

Related Links

- security key-manager external boot-interfaces modify

security key-manager key show

Display Encryption Key IDs

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

Description

This command displays the key IDs of the authentication keys (NSE-AK) and vserver keys (SVM-KEK) that are available in onboard key management. This command is not supported for an external key management 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.

| [-detail]

If this parameter is specified, the command displays additional details about the key IDs.

| [-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 key IDs that are located on the specified storage system.

[-key-store <Key Store>] - Key Store

If this parameter is specified, the command displays information only about key IDs that are managed by the specified key management. For example, use *onboard* for onboard key management.

[-key-id <text>] - Key Identifier

If this parameter is specified, the command displays information only about the specified key IDs.

[-key-tag <text>] - Key Tag

If this parameter is specified, the command displays information only about key IDs that have the specified key tags.

[-key-location <text>] - Key Location

If this parameter is specified, the command displays information only about key IDs that are located on the specified key location. For example, use *local-cluster* for onboard key management.

[-used-by <Key Usage Type>] - Used By

If this parameter is specified, the command displays information only about key IDs that are associated with the specified application usage of the keys. For example, "NSE-AK" would display key IDs only for NSE drives.

[-restored {yes|no}] - Restored

If this parameter is specified, the command displays information only about key IDs that have the specified value of restored keys. If restored is *yes* , then the corresponding key is available (normal). If restored is *no* , use the security key-manager setup command to restore the key. See the man page for security key-manager setup for details.

Examples

The following example shows all keys stored in the onboard key manager:

```
cluster-1::> security key-manager key show

Node: node1
Key Store: onboard
Key ID                                                                   Used By
------------------------------------------------------------------------ --------
000000000000000002000000000001001BC4C708E2A89A312E14B6CE6D4D49D4 NSE-AK
000000000000000002000000000001005E89099721F8817E65E3AEB68BE1BFCA NSE-AK
000000000000000002000000000000A0046DF92864D4CECE662B93BEB7F536610 SVM-KEK

Node: node2
Key Store: onboard
Key ID                                                                   Used By
------------------------------------------------------------------------ --------
000000000000000002000000000001001BC4C708E2A89A312E14B6CE6D4D49D4 NSE-AK
000000000000000002000000000001005E89099721F8817E65E3AEB68BE1BFCA NSE-AK
000000000000000002000000000000A0046DF92864D4CECE662B93BEB7F536610 SVM-KEK
6 entries were displayed.
```

The following example shows a detailed view of all keys stored in the onboard key manager:

```
cluster-1::> security key-manager key show -detail

Node: node1
Key Store: onboard
Key ID Key Tag         Used By    Stored In
Restored
------ --------------- ---------- ----------------------------------
--------
000000000000000020000000000001001BC4C708E2A89A312E14B6CE6D4D49D4
       -               NSE-AK     local-cluster                      yes
000000000000000020000000000001005E89099721F8817E65E3AEB68BE1BFCA
       -               NSE-AK     local-cluster                      yes
000000000000000020000000000A0046DF92864D4CECE662B93BEB7F536610
       -               SVM-KEK    local-cluster                      yes

Node: node2
Key Store: onboard
Key ID Key Tag         Used By    Stored In
Restored
------ --------------- ---------- ----------------------------------
--------
000000000000000020000000000001001BC4C708E2A89A312E14B6CE6D4D49D4
       -               NSE-AK     local-cluster                      yes
000000000000000020000000000001005E89099721F8817E65E3AEB68BE1BFCA
       -               NSE-AK     local-cluster                      yes
000000000000000020000000000A0046DF92864D4CECE662B93BEB7F536610
       -               SVM-KEK    local-cluster                      yes
6 entries were displayed.
```

Related Links

- security key-manager setup

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 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 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}] - Second Authentication Method2

This specifies the authentication method for the login. It will be used as the second factor for authentication. Possible values include the following:

- password - Password
- publickey - Public-key authentication
- none - default value

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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* . Here *is-ns-switch-group* must be set to *yes* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -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 *vs* :

```
cluster1::> security login create -vserver vs -user-or-group-name monitor
-application ssh -authentication-method password -second-authentication
-method none -role vsadmin
```

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 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 *vs* :

```
cluster1::> security login delete -user-or-group-name guest -application
ssh -authentication-method password -vserver vs
```

The following example illustrates how to delete a login that has the username *guest* , the application *ontapi* , and the authentication method *cert* for Vserver *vs* :

```
cluster1::> security login delete -user-or-group-name guest -application
ontapi -authentication-method cert -vserver vs
```

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 *vs* :

```
cluster1::> security login delete -user-or-group-name DOMAIN1\adgroup
-application ssh -authentication-method domain -vserver vs
```

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 *vs* :

```
cluster1::> security login delete -user-or-group-name nssgroup
-application ontapi -authentication-method nsswitch -vserver vs
```

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 alogrithm using `-hash -function` parameter.

Examples

The following command expires the password of the 'jdoe' user account which belongs to the 'vs1' Vserver.

```
cluster1::> security login expire-password -vserver vs1 -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' Vserver user account named 'jdoe' that is encrypted with the MD5 hash function.

```
cluster1::> security login expire-password -vserver vs1 -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.

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'.

```
cluster1::> security login lock -vserver vs1 -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 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}] - Second Authentication Method2

This specifies the authentication method for the login method. It will be used as the second factor for authentication. Possible values include the following:

- password - Password
- publickey - Public-key authentication
- none - default value

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 *vs* :

```
cluster1::> security login modify -user-or-group-name guest -application
ontapi -authentication-method password -role guest -vserver vs
```

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 *vs* :

```
cluster1::> security login modify -user-or-group-name guest -application
ssh -authentication-method publickey -role vsadmin -vserver vs
```

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 *vs* . 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 vs -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 *vs* :

```
cluster1::> security login modify -user-or-group-name guest -application
ssh -authentication-method publickey -second-authentication-method
password -vserver vs
```

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 *vs* :

```
cluster1::> security login modify -user-or-group-name guest -application
ssh -authentication-method publickey -second-authentication-method none
-vserver vs
```

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 Data 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 Data 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 Data 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 'vs' Vserver.

```
cluster1::> security login password -username admin -vserver vs
```

The following command initiates a password change for the 'vs' Vserver user account named 'admin'. The new password will be encrypted by using the SHA512 password-hashing algorithm.

```
cluster1::*> security login password -username admin -vserver vs -hash
-function sha512
```

The following command initiates a password change for the 'vs' Vserver user account named 'admin'. The new password will be encrypted by using the SHA256 password-hashing encryption algorithm.

```
cluster1::*> security login password -username admin -vserver vs -hash
-function sha256
```

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 (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

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 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}] - Second Authentication Method2

Displays the login methods that match the specified authentication method to be used as the second factor. Possible values include the following:

- password - Password
- publickey - Public-key authentication
- none - default value

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          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

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
9 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.

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.

```
cluster1::> security login unlock -vserver vs1 -username jdoe
```

security login whoami

Show the current 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 and role 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
```

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 {(ftp|http)://(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, Data 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 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 {(ftp|http)://(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.

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.

`[-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.

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”.

```
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"
```

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.

-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.

```
cluster1::> security login publickey delete -username tsmith -index 5
```

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 {(ftp|http)://(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.

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.

`-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.

Examples

The following command modifies the public key at index number 10 for the 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"
```

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.

[-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.

Examples

The example below displays public key information for the user named tsmith.

```
cluster1::> security login publickey show -username tsmith
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
```

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. To specify the default setting, use the special value "`DEFAULT` ".

`[-access <Access>]` - 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 <Access>]` - 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 Data 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 Data 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 Data 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 Data 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 Data 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

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 <Access>] - 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 <unsigned32_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 *0* , which means failed login attempts do not cause an account to be locked.

[-lockout-duration <integer>] - Maximum Lockout Period (Days)

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 day.

[-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 charaters. The default setting is *0* , which requires no digits.

[-passwd-expiry-warn-time <unsigned32_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 <unsigned32_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 <unsigned32_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.

Examples

The following command modifies the user-account restrictions for an account with the role name admin for a Vserver named vs. The minimum size of the password is set to 12 characters.

```
cluster1::> security login role config modify -role admin -vserver vs
-passwd-minlength 12
```

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 Data 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")
- 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 ("0")
- Number of days that the user account is locked out after the maximum number of failed login attempts is reached ("0")

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`
- Number of days for which the user account is locked after the maximum number of failed login attempts is reached `-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`

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 <unsigned32_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>] - 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 <unsigned32_or_unlimited>] - Display Warning Message Days Prior to Password Expiry (Days)

Selects the profile configurations that match this parameter value.

[-account-expiry-time <unsigned32_or_unlimited>] - Account Expires in (Days)

Selects the profile configurations that match this parameter value.

[-account-inactive-limit <unsigned32_or_unlimited>] - Maximum Duration of Inactivity before Account Expiration (Days)

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 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] }

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

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

Displays the insecure 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 {(ftp|http)://(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 Data 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 Data 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 necesary 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 paramater 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 contaning 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] }

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 an active 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 active 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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 active CLI & ONTAPI 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}` - Interface

The interface (CLI or ONTAPI) 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}` - Interface
The interface (CLI or ONTAPI) 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}` - Interface
The interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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] }

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

[-interface {cli|ontapi}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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] }

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

[-interface {cli|ontapi}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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] }

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

[-interface {cli|ontapi}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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} - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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} - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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] }

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

[-interface {cli|ontapi}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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}` - Interface (privilege: advanced)

The interface (CLI or ONTAPI) 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 ] }`

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

[-interface {cli|ontapi}] - Interface (privilege: advanced)

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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}] - Interface

Selects the sessions that match this parameter value. This identifies the interface (CLI or ONTAPI) 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. Data ONTAP supports the *diffie-hellman-group-exchange-sha256* key exchange algorithm for SHA-2. Data 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. Data ONTAP also supports *ecdh-sha2-nistp256* , *ecdh-sha2-nistp384* , *ecdh-sha2-nistp521* , and *curve25519-sha256* . Data 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* . Data ONTAP supports MAC algorithms of the following types: *hmac-sha1* , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *hmac-ripemd160* , *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* , *hmac-ripemd160-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.

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
```

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 or ciphers or MAC 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. Data ONTAP supports the *diffie-hellman-group-exchange-sha256* key exchange algorithm for SHA-2. Data 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. Data 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* . Data ONTAP supports MAC algorithms of the following types: *hmac-sha1* , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *hmac-ripemd160* , *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* , *hmac-ripemd160-etm* , *umac-64-etm* , and *umac-128-etm* .

Parameters

-vserver <Vserver Name> - Vserver

Identifies the Vserver for which you want to replace the existing SSH key exchange algorithm and cipher 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.

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 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
-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 Data 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 Data 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 is 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 will notdoes not fail. Data ONTAP supports the *diffie-hellman-group-exchange-sha256* key exchange algorithm for SHA-2. Data 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. Data ONTAP also supports *ecdh-sha2-nistp256* , *ecdh-sha2-nistp384* , *ecdh-sha2-nistp521* , and *curve25519-sha256* . Data 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* . Data ONTAP supports MAC algorithms of the following types: *hmac-sha1* , *hmac-sha1-96* , *hmac-md5* , *hmac-md5-96* , *hmac-ripemd160* , *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* , *hmac-ripemd160-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.

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 and maximum
authentication retry count 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. Data ONTAP
supports the ``_diffie-hellman-group-exchange-sha256_`` key exchange
algorithm for SHA-2. Data ONTAP also supports the ``_diffie-hellman-group-
exchange-sha1_`` , ``_diffie-hellman-group14-sha1_`` , and ``_diffie-
hellman-group1-sha1_`` key exchange algorithms for SHA-1. Data ONTAP also
supports ``_ecdh-sha2-nistp256_`` , ``_ecdh-sha2-nistp384_`` , ``_ecdh-
sha2-nistp521_`` , ``_curve25519-sha256_`` . Data 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_`` . Data ONTAP supports MAC algorithms
of the following types: ``_hmac-sha1_`` , ``_hmac-sha1-96_`` , ``_hmac-
md5_`` , ``_hmac-md5-96_`` , ``_hmac-ripemd160_`` , ``_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_`` ,
``_hmac-ripemd160-etm_`` , ``_umac-64-etm_`` , ``_umac-128-etm_``
```

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 algorithm, cipher, and MAC algorithm configurations.

[-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.

Examples

The following command displays the enabled SSH key exchange algorithms, ciphers, MAC 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:

```
cluster-1::> security ssh show
                            Key Exchange  MAC              Max
Authentication
Vserver        Ciphers     Algorithms    Algorithms              Retry
Count
-------------- ----------  ------------  --------------
------------------
cluster-1      3des-cbc    diffie-       hmac-sha1
4
                           hellman-
                           group-
                           exchange-
                           sha256
vs1            aes256-     diffie-       hmac-sha1,
6
               ctr,        hellman-      hmac-sha1-96,
               aes192-     group-        hmac-sha2-256,
               ctr,        exchange-     hmac-sha2-512,
               aes128-     sha256,       hmac-sha1-etm,
               ctr,        diffie-       hmac-sha1-96-
               aes256-     hellman-      etm,
               cbc,        group-        hmac-sha2-256-
               aes192-     exchange-     etm,
               cbc,        sha1,         hmac-sha2-512-
               aes128-     diffie-       etm, hmac-md5,
               cbc,        hellman-      hmac-md5-96,
               3des-cbc,   group14-      hmac-
               aes128-     sha1,         ripemd160,
               gcm,        ecdh-sha2-    umac-64,
               aes256-gcm  nistp256,     umac-128,
                           ecdh-sha2-    hmac-md5-etm,
                           nistp384,     hmac-md5-96-
                           ecdh-sha2-    etm,
                           nistp521,     hmac-
                           curve25519-   ripemd160-etm,
                           sha256        umac-64-etm,
                                         umac-128-etm
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 occured. 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>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 <unsigned32_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>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>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 <unsigned32_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>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

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. System ComplianceClock can be initialized only once by the user. Once initialized, user cannot make any changes to the System ComplianceClock. Hence, user should ensure that system clock is correct before initializing the System ComplianceClock.

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 paramter, 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* 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 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.

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 begin 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 modfied. 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 Vserveers 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, 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.

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 or Infinite Volume relationships.

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>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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 copy 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 copy 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:
```

Related Links

- snapmirror show

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 or SolidFire endpoint. When Data ONTAP breaks the relationship, if the endpoint is a volume 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 copy. 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.

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 Data ONTAP revert operation (see the `-delete-snapshots` parameter in advanced privilege level). This command is not intended for use with load-sharing mirrors.

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>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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 Snapshot copies on a vault destination so that the system can be reverted. Note that the only Snapshot copies that will be deleted are those that were created with the current version of Data ONTAP. Any Snapshot copies that might be present created with a different version will not be deleted.

[-s, -restore-destination-to-snapshot <text>] - Restore Destination to Snapshot Copy

This optional parameter specifies the Snapshot copy 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 copy. This parameter is not supported for Vserver or FlexGroup relationships.

[-recover <true>] - Recover (privilege: advanced)

When a SnapMirror break operation fails on a FlexGroup relationship, a subset of the destination FlexGroup 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 copy, and any writes to the read-write constituents will be lost. This parameter is applicable only for SnapMirror relationships with FlexGroup 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:
```

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 Infinite Volumes and between Vservers. Infinite Volumes and Vservers support only data protection relationships. A SnapMirror relationship between Vservers can only be created if the system containing the source Vserver is also running Data ONTAP 8.3 or later. Creating Vserver SnapMirror relationships is not supported for Vservers with Infinite Volumes. You can also use the `snapmirror create` command to create an extended data protection relationship between FlexGroups. FlexGroups only support extended data protection relationships. A SnapMirror relationship between FlexGroups is only supported if the system containing the source FlexGroup volume is also running Data ONTAP 9.1.0 or later. The source or destination of a FlexGroup SnapMirror relationship cannot be the source or destination of any other SnapMirror relationship.

The `snapmirror create` command can be used to create an extended data protection (XDP) relationship between a Data ONTAP volume and a non-Data ONTAP endpoint that supports SnapMirror (AltaVault, Solidfire). AltaVault endpoints can only be used as destinations. SolidFire endpoints can be used as sources or destinations.

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 FlexGroups 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.

For an Infinite Volume SnapMirror relationship, the destination Infinite Volume size must be greater than or equal to the source Infinite Volume size in bytes. You can verify the size in bytes by running set`-units `*`KB`* followed by a volume show command.

If all systems involved are running Data 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 Data ONTAP 8.1, if the source volume is on a storage system running clustered Data 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 Data 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 Data 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 Data 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. Load sharing mirrors are not supported for Infinite Volumes.

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>} - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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, an extended data protection (XDP) relationship, a transition data protection relationship between a Data ONTAP operating in 7-Mode system and a clustered Data ONTAP system (TDP), or a load-sharing (LS) relationship. The default value is XDP for all relationships except for Infinite Volume and Vserver DR relationships. Infinite Volumes support only DP relationships. FlexGroup volumes and non-Data 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, 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.

The schedules associated with an Infinite Volume SnapMirror relationship should not have an interval shorter than hourly.

[-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 no type is specified or when the data protection (DP) type is specified is *MirrorAllSnapshots* and the default policy when the extended data protection (XDP) type is specified is *XDPDefault* . For FlexGroup volume relationships, the *MirrorAllSnapshots* policy 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 clustered Data ONTAP 8.2 data protection (DP) relationships were used for mirroring, while extended data protection (XDP) relationships were used for vaulting. In clustered Data 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 clustered Data 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 Data ONTAP source volume and an AltaVault destination endpoint, only policies of type *vault* are supported. For XDP relationships between a Data 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 Data 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 <unsigned32_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 which 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.

[-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.

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 *MirrorAllSnapshots*, 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
    -type DP
```

To create an extended data protection relationship between the source FlexGroup volume *vs1.example.com:fg_src* and the destination FlexGroup volume *vs2.example.com:fg_dst*, with the default policy of *MirrorAllSnapshots*, 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 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 Data 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/mkt1* , and the destination endpoint *cluster1://vs1.example.com/mkt1_ls1* with the schedule named *5min* used to update the relationship, type the following command:

```
cluster1::> snapmirror create
       -destination-path cluster1://vs1.example.com/mkt1_ls1
       -source-path      cluster1://vs1.example.com/mkt1 -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 Data 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 Data 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
```

Related Links

- vserver create
- vserver peer create
- snapmirror initialize
- volume create

- snapmirror show
- snapmirror list-destinations
- set
- volume show
- 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, or a non-Data ONTAP endpoint. The Vservers, volumes and non-Data ONTAP destinations are not destroyed and Snapshot copies 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, Snapshot copies that are locked by SnapMirror on the source container cannot be cleaned up using the snapmirror release command. If the source container is a Data ONTAP volume, in order to reclaim space captured in the base Snapshot copy on the volume, issue a `snapshot delete` command specifying the `-ignore-owners` parameter in diag privilege level. To reclaim space captured in a Snapshot copy 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.

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. The use of `snapmirror delete` on a source cluster is not supported for an Infinite Volume relationships in this release.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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:
```

Related Links

- snapmirror release
- snapmirror list-destinations
- snapmirror show

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 Data ONTAP takes a Snapshot copy on the source volume to capture the current image of the source volume and transfers all of the Snapshot copies on the source volume to each of the destination volumes.

After the `snapmirror initialize-ls-set` command successfully completes, the last Snapshot copy transferred is made the exported Snapshot copy on the destination volumes.

The parameter that identifies the set of load-sharing mirrors is the source volume. Data and Snapshot copies 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 not supported on Infinite Volume 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.

Parameters

{ -S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 the 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-Data 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-Data 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/dept_eng* , type the following command:

```
cluster1::> snapmirror initialize-ls-set -source-path
       //vs1.example.com/dept_eng
```

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 or a non-Data ONTAP endpoint of a SnapMirror relationship. The command behaves differently between data protection (DP), extended data protection (XDP) and load-sharing (LS) relationships.

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 copy 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 Snapshot copies up to and including the Snapshot copy 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 Snapshot copies up to and including the Snapshot copy created by it or only the Snapshot copy 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 copy created by it.

After the `snapmirror initialize` command successfully completes, the last Snapshot copy transferred is made the exported Snapshot copy on the destination volume.

When an Infinite Volume SnapMirror relationship is initialized, the command will create any needed constituent volumes for the destination Infinite Volume. The Infinite Volume relationship will appear in the snapmirror show command output on the source cluster after it is initialized.

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 Infinite Volumes and 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>}] - 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 the 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-Data 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-Data ONTAP

source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For relationships with 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-Data 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-Data ONTAP destination endpoints.

[-s, -source-snapshot <text>] - Source Snapshot

This optional parameter specifies the Snapshot copy that `snapmirror initialize` will use for the baseline transfer. For data protection (DP) relationships, the baseline transfer will include all of the Snapshot copies up to and including the specified Snapshot copy. 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 Snapshot copies up to and including the specified Snapshot copy or only the specified Snapshot copy 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 copy. This parameter is not supported for relationships with *"Relationship Capability"* of *"Pre 8.2"*. This parameter is not supported for Infinite Volume SnapMirror relationships.

[-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"*. This parameter is not supported by this operation for Infinite Volumes.

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. 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"*.

[-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.

[-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 Data 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 Data 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:
```

Related Links

- snapmirror show
- snapmirror initialize-ls-set
- snapmirror create
- 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 copy 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-Data 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>}] - 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>}] - 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|infinitevol|consistencygroup|flexgroup}] - Relationship Group Type

Selects SnapMirror relationships that have a matching relationship group type. Possible values are:

- none
- vserver
- infinitevol
- flexgroup

[-policy-type {vault|async-mirror|mirror-vault}] - SnapMirror Policy Type

Selects SnapMirror relationships that have a matching SnapMirror policy type. Possible values are:

- async-mirror
- vault
- mirror-vault

[-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 SnapMirror relationships, but it is supported for FlexGroup 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 Infinite Volume or FlexGroup SnapMirror relationships, but it is supported for FlexGroup 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 Infinite Volume or FlexGroup SnapMirror relationships, but it is supported for FlexGroup constituent relationships.

[-is-constituent {true|false}] - Constituent Relationship

Selects SnapMirror relationships that have a matching constituent condition. This parameter is not supported for Vserver, FlexGroup, or FlexGroup constituent SnapMirror 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.
```

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, or a non-Data 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>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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.

The schedules associated with an Infinite Volume SnapMirror relationship should not have an interval shorter than hourly.

[-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.

[-tries <unsigned32_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 which 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.

[-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
```

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 copy, 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"*.

This command is not supported on Infinite Volume snapmirror relationships.

Parameters

{ [-S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For relationships with 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-Data 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-Data 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 *dept_eng_ls_mirror1* to be the source read-write volume for mirroring and client access, type the following command:

```
cluster1::> snapmirror promote -destination-path
         //vs1.example.com/dept_eng_ls_mirror1
         -source-path //vs1.example.com/dept_eng -f true
```

Related Links

- snapmirror show

snapmirror protect

Start protection for 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 a list of volumes. For each volume, the command creates a data protection destination volume 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 Infinite Volumes or Infinite Volume constituents, FlexGroup volumes or FlexGroup volume constituents, Vserver endpoints or non ONTAP endpoints.

Parameters

-path-list {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - Path List

This parameter specifies the list of volumes to be protected. The list is a comma separated list of paths of the form *vserver:volume*, for example *vs1.example.com:dept_eng1, vs1.example.com:dept_eng2*

[-destination-vserver <vserver name>] - Destination Vserver

This parameter specifies the Vserver in which to create the destination volumes of the SnapMirror relationships.

[-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.

[-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.

[-destination-volume-suffix <text>] - Destination Volume Name Suffix

This optional parameter designates the suffix for the destination volume name. If you do not desginate 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.

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
```

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.

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.

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>}] - 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 the 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-Data 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-Data ONTAP

source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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:
```

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 copy owner tags and any Snapshot copies 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.

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-Data 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>}] - Source Path

Specifies the source endpoint of the SnapMirror relationship in one of two 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.

| [-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>} - Destination Path

Specifies the destination endpoint of the SnapMirror relationship in one of two 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.

| -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 Snapshot copies 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:
```

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 copy or one or more files or LUNs of a Snapshot copy from one volume to another volume.

The source of the restore can be a volume that is:

- the destination volume 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 which is not the destination endpoint of any SnapMirror relationship
- a read-write volume.
- an AltaVault endpoint. In this case the destination must be an empty Data ONTAP volume.

The following cannot be used as either the source or destination volume of a restore:

- a volume 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.
- an Infinite Volume.

A SnapMirror relationship of type *RST* is created from the source volume to the destination volume 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 copy to a destination volume.

By default the `snapmirror restore` will copy the latest Snapshot copy from the source volume to the destination volume. A specific Snapshot copy 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 copy. 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 copy have been restored.

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 copy selected to be restored are copied to the active file system of the destination volume.
- The destination volume 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 copy between the source and destination volumes. Restoring a Snapshot copy to an empty read-write volume is not supported. Incremental restore from a non-Data ONTAP endpoint to a Data ONTAP volume is not supported.

An incremental restore preserves all Snapshot copies on the destination volume but does not preserve changes to the active file system since the latest Snapshot copy. To preserve changes to the destination volume since the latest Snapshot copy 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 copy on the destination volume and the destination volume is made read-only.
- This Snapshot copy is the exported Snapshot copy and it is the view to which clients are redirected when accessing the destination volume.
- The contents of the Snapshot copy selected to be restored are copied to the active file system of the destination volume.
- The destination volume 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 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 copy has not been restored to the destination volume, all data copied to the active file system by `snapmirror restore` to the destination volume is reverted.
- The destination volume 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 copy has been restored to the destination volume, the volume is made read-write.
- The *RST* SnapMirror relationship is deleted.

The following paragraphs describe the behavior and requirements when restoring one or more files or LUNs to the destination volume.

The destination volume must be a read-write volume. Restoring files or LUNs to a data protection volume is not supported. When restoring files or LUNs the source and destination volumes are not required to have a common Snapshot copy. If a common Snapshot copy exists, an incremental restore is performed for those files or LUNs being restored which exist in the common Snapshot copy.

The contents of the files or LUNs to which data is being restored on the destination volume are not preserved by this command. To preserve the contents of the destination files or LUNs, create a Snapshot copy on the destination volume prior to running this command. Client I/O is not allowed to a file or LUN to which data is being restored on the destination volume.

The `-source-snapshot` parameter is required when restoring files or LUNs. It identifies the Snapshot copy on the source volume from which the files or LUNs to be restored are copied. If all files or LUNs to be restored do not exist in this Snapshot copy the command fails.

The source path for each file or LUN being restored is required. The source path of a file or LUN is from the root of the source Snapshot copy 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 or LUN to which data is being restored on the destination volume does not exist, the file or LUN is created. If any directory in the path of the file or LUN 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 or overwriting the contents of an existing LUN with the contents of a different LUN is supported. Overwriting a file with a LUN or a LUN with a file is not supported. Client I/O is not allowed to all files and LUNs 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 or LUNs to the same destination volume, are not supported. The destination volume of a `snapmirror restore` to which one or more file or LUNs are being restored, can simultaneously be the source volume of a snapmirror update .

For a file or LUN restore the following steps are performed:

- Create the *RST* SnapMirror relationship.
- If any file or LUN being restored does not exist on the destination volume, create all such files or LUNs.
- Prevent client I/O to files or LUNs 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 or LUNs to the corresponding file or LUN on the destination volume.
- Allow client I/O to files or LUNs to which data has been restored on the destination volume.
- Delete the *RST* SnapMirror relationship.

Some file restore operations require a Snapshot copy to be created. This Snapshot copy is temporary, it is deleted before the operation completes.

Since client I/O is not allowed to files or LUNs being restored, client I/O to files or LUNs being restored should be quiesced. Mapped LUNs remain mapped throughout the operation. SAN clients do not need to rediscover a mapped LUN that has been restored.

If `snapmirror restore` fails or is aborted, the *RST* relationship remains. Use the snapmirror show command with the destination volume 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 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 copy 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.

The `snapmirror restore` 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>}] - Source Path

Specifies the source endpoint in one of three 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. A non-Data 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-Data 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-Data 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-Data ONTAP endpoint.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - Destination Path

Specifies the destination endpoint in one of two 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-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 copy, this optional parameter identifies the Snapshot copy 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 or LUNs from a Snapshot copy, 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` .

[-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 copy 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 or LUNs of a Snapshot copy, any file to which client I/O is not allowed is deleted.

[-tries <unsigned32_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 or LUNs to be restored. The list can contain specifications for up to 8 files or LUNs. Specification for each file or LUN 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 copy, 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. 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
```

[-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-Data ONTAP endpoints.

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 copy 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 restores /file3 from the source Snapshot copy *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 copy *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 copy *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

- volume quota modify
- volume snapshot create
- snapmirror show
- snapmirror update
- lun delete

snapmirror resume

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.

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 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>|<hostip:/share/share-name>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data 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:
```

Related Links

- snapmirror show

snapmirror resync

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 or non-Data 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.
- snapmirror update command failed because the required common Snapshot copy 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 copy. In this case, `snapmirror resync` might implicitly create the SnapMirror relationship between them.

The `snapmirror resync` command can cause data loss on the destination volume because the command can remove the exported Snapshot copy 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 copy between the source and destination volumes, removes Snapshot copies on the destination volume that are newer than the common Snapshot copy and mounts the destination volume as a DP volume with the common Snapshot copy as the exported Snapshot copy.
- For data protection (DP) relationships, takes a Snapshot copy of the source volume to capture the current image and transfers Snapshot copies that are newer than the common Snapshot copy from the source volume to the destination volume. For extended data protection (XDP) relationships, transfers Snapshot copies newer than the common Snapshot copy according to the relationship policy, i.e., Snapshot copies

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 copy of the source volume and includes it in the Snapshot copies selected for transfer.

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* .

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-Data ONTAP endpoint. When used, the parameter *"preserve"* changes the behavior of `snapmirror resync` command. Changed behavior of the command can be described as follows:

- Finds the most recent common Snapshot copy between the source and destination volumes, preserves all Snapshot copies on the destination volume that are newer than the common Snapshot copy, and mounts the destination volume as a DP volume with the common Snapshot copy as the exported Snapshot copy.
- Performs a local rollback transfer to make a copy of the common Snapshot copy on the destination volume and establish it as the latest Snapshot copy on the destination volume. The command then transfers all Snapshot copies that are newer than the common Snapshot copy, from the source volume to the destination volume. The command only transfers Snapshot copies that match the relationship's policy, i.e., Snapshot copies 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 snapmirror create command, the `snapmirror resync` command will implicitly create the SnapMirror relationship, with the same behaviors as described for the snapmirror create command before resyncing it.

For Infinite Volumes and Vservers, you must create SnapMirror relationships between Infinite Volumes or Vservers by using the 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 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 resync` command fails if the destination volume does not have a Snapshot copy 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>|<hostip:/share/share-name>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name* . For relationships with 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-Data 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-Data ONTAP destination endpoints.

[-s, -source-snapshot <text>] - Source Snapshot

This optional parameter specifies a Snapshot copy to transfer. The default behavior for a data protection relationship with a read-write source is that Data ONTAP creates a new Snapshot copy and uses it as the basis for determining what data are replicated; with this option, the specified Snapshot copy 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 copy must be newer than the latest common Snapshot copy. For an extended data protection relationship, the specified Snapshot copy can be newer or older than the common Snapshot copy. 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 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"*. This parameter is not supported by this operation for Infinite Volumes.

You define and name a policy using the 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 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. 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-Data ONTAP endpoints. When specified, it changes the behavior of the `snapmirror resync` command to preserve Snapshot copies on the destination volume that are newer than the latest common Snapshot copy. This parameter is not supported for relationships with *"Relationship Capability"* of *"Pre 8.2"* . This parameter is not supported for Infinite Volume SnapMirror relationships.

[-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-Data ONTAP endpoints. 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.

[-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.

[-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 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 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 Data 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 Data ONTAP source endpoint *vs1.example.com:data_ontap_vol* , and the non-Data 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 snapmirror break command, type the following command:

```
cluster2::> snapmirror resync -destination-path
         dvs1.example.com:
```

Related Links

- snapmirror update
- snapmirror create
- snapmirror show
- job show

- job history show
- snapmirror policy create
- snapmirror modify
- snapmirror break

snapmirror set-options

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. This command is not supported for Infinite Volume SnapMirror relationships.

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
```

snapmirror show-history

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 history of the SnapMirror operations. It is not supported for FlexGroup SnapMirror relationships. This command is also not supported for relationships with non-Data 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 about operations which happened prior to installing Data ONTAP 8.3. Also, it does not return information for relationships with the *"Relationship Capability"* field, as shown in the output of the snapmirror show command, set to *"Pre 8.2"*.

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.
                              - infinitevol: Infinite Volume
                                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>|<hostip:/share/share-name>}] - 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>|<hostip:/share/share-name>}] - 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}] - 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|infinitevol|consistencygroup|flexgroup}] - Relationship Group Type

Select SnapMirror relationships that have a matching relationship group type. Possible values are:

- none
- vserver
- infinitevol

[-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 a matching additional information.

[-max-rows-per-relationship <integer>] - Maximum Number of Rows per Relationship

Select 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.
```

snapmirror show

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 clustered Data ONTAP 8.1, and to accommodate load-sharing relationships which are only supported in a Data ONTAP 8.1 compatible way, SnapMirror relationships, which match one of the following conditions are managed as on clustered Data ONTAP 8.1: (1) The relationship is of type load-sharing; (2) The source endpoint of the relationship is on a remote Data ONTAP 8.1 cluster; (3) The local cluster was upgraded from clustered Data ONTAP 8.1, the relationship was created before the upgrade, and the relationship has not yet been converted to one with Data ONTAP 8.2 capabilities. These relationships have the same limitations as on clustered Data 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-Data ONTAP endpoint that supports SnapMirror (for example, AltaVault). For backward compatibility with clustered Data 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 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 clustered Data 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.
                             - infinitevol: Infinite Volume relationship.
                             - flexgroup: FlexGroup 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.
```

```
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.
                    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 copy on the
                            destination volume.
    Newest Snapshot Timestamp: Timestamp of the newest Snapshot copy.
            Exported Snapshot: Name of the exported Snapshot copy on the
                            destination volume.
  Exported Snapshot Timestamp: Timestamp of the exported Snapshot copy.
                     Lag Time: Time since the exported Snapshot copy
                            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 copy 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. This parameter
                              is not supported for SnapMirror FlexGroup
                              relationships, but it is supported for
                              FlexGroup constituent relationships.
    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
                              clustered Data ONTAP 8.1.
                            - "8.2 and above": Full support of clustered
                              Data 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
                            xref:{relative_path}snapmirror-policy-
create.html[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.
    Constituent Relationship: Whether or not the SnapMirror relationship
                            is between Infinite Volume constituent
                            volumes. Can be:
                            - true:  The relationship is between
                                     constituent volumes.
                            - false: The relationship is not between
                                     constituent volumes.
                            This parameter is not supported for
                            FlexGroup or FlexGroup constituent
                            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.
                            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.
Number of Successful Updates: The number of successful SnapMirror update
                              operations for 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.
                              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.
                              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.
                              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.
                               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.
                               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
                               or FlexGroup SnapMirror relationships, but
                               it is supported for FlexGroup constituent
                               relationships.
          Total Transfer Time: Cumulative total transfer time in seconds
                               for 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>|<hostip:/share/share-name>}] - 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-Data ONTAP 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-Data 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-Data ONTAP source endpoints.

{ [-destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>}] - Destination Path

Select SnapMirror relationships that have a matching destination path name.

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-Data ONTAP 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-Data 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-Data ONTAP destination endpoints.

[-type <snapmirrorType>] - Relationship Type

Select SnapMirror relationships that have a matching relationship type. Infinite Volumes and Vservers support only DP SnapMirror relationships. Possible values are:

- DP
- LS
- XDP
- TDP
- RST

[-relationship-group-type {none|vserver|infinitevol|consistencygroup|flexgroup}] - Relationship Group Type

Select SnapMirror relationships that have a matching relationship group type. Possible values are:

- none
- vserver
- infinitevol
- flexgroup

[-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}] - SnapMirror Policy Type

Selects SnapMirror relationships that have a matching SnapMirror policy type. Possible values are:

- async-mirror
- vault
- mirror-vault

[-policy <sm_policy>] - SnapMirror Policy

Select SnapMirror relationships that have a matching SnapMirror policy.

[-tries <unsigned32_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.

[-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 and Waiting are not supported for Infinite Volume SnapMirror relationships. 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 copy.

[-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 copy checkpoint. This parameter is not supported for Vserver SnapMirror relationships.

[-newest-snapshot <text>] - Newest Snapshot

Select SnapMirror relationships that have a matching newest Snapshot copy.

[-newest-snapshot-timestamp <MM/DD HH:MM:SS>] - Newest Snapshot Timestamp

Select SnapMirror relationships that have a matching newest Snapshot copy timestamp.

[-exported-snapshot <text>] - Exported Snapshot

Select SnapMirror relationships that have a matching exported Snapshot copy 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 copy 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|cggrs_resync|cg_update|cg_initialize|cg_restore}] - Transfer Type

Select SnapMirror relationships that have a matching current transfer type. Transfer type Check is not supported for Infinite Volume SnapMirror relationships.

[-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|cggrs_resync|cg_update|cg_initialize|cg_restore}] - 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. This option is not supported for Infinite Volume SnapMirror relationships.

[-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. This option is not supported for Infinite Volume SnapMirror relationships.

[-is-constituent {true|false}] - Constituent Relationship

Select SnapMirror relationships that have a matching constituent condition.

[-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 and FlexGroup 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

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: -
              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: -
               Constituent Relationship: false
                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: -
           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: -
              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: -
               Constituent Relationship: false
                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: -
```

```
           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: -
              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: -
               Constituent Relationship: false
                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: -
            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: -
              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: -
                  Constituent Relationship: false
                   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: -
            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: -
        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: -
         Constituent Relationship: false
    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: -
     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.
```

Related Links

- snapmirror list-destinations
- snapmirror policy create

snapmirror update-ls-set

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, Data ONTAP takes a Snapshot copy on the source volume to capture the current image of the source volume, finds the most recent common Snapshot copy between the source and destination volumes, and incrementally transfers Snapshot copies that are newer than the common Snapshot copy to the destination volume.

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 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 copy transferred is made the new exported Snapshot copy on the destination volumes.

This command is not supported on Infinite Volume 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.

Parameters

{ -S, -source-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 the 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-Data 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-Data 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/dept_eng*, type the following command:

```
cluster1::> snapmirror update-ls-set -source-path
//vs1.example.com/dept_eng
```

Related Links

- snapmirror update
- snapmirror show

snapmirror update

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 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 snapmirror initialize or 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 copy on the source volume to capture the current image of the source volume

- Finds the most recent Snapshot copy on the destination volume and validates that the corresponding Snapshot copy is still present on the source
- Incrementally transfers Snapshot copies that are newer than the corresponding Snapshot copy 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.

Using the `snapmirror update` command to update a set of load-sharing mirrors will not work. Use the 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 copy on the source volume. Depending on the rules in the policy, the command might transfer just the newly created Snapshot copy or all Snapshot copies that are newer than the common Snapshot copy including the newly created Snapshot copy 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 copy on the source volume but transfers only selected Snapshot copies that are newer than the common Snapshot copy to the destination volume. (Those older than the common copy can be transferred by using the `-source-snapshot` parameter.) Snapshot copies are selected by matching the value of `-snapmirror-label` of a Snapshot copy with the value of `-snapmirror-label` of one of the rules from the corresponding SnapMirror policy associated with the SnapMirror relationship. All matching Snapshot copies 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 copy on the source volume and transfers only selected Snapshot copies that are newer than the common snapshot copy. The newly created Snapshot copy 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 Snapshot copies on the destination volume. It does so by deleting Snapshot copies that have exceeded the value of `-keep` for the matching rule from the corresponding SnapMirror policy associated with the SnapMirror relationship. Snapshot copies that match the same `-snapmirror-label` will be deleted in oldest-first order.

For data protection relationships, the parameter `-source-snapshot` is optional and only allows for the transfer of Snapshot copies newer than the common Snapshot copy 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 copy that is older than the common Snapshot copy 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 copy on the source volume. Depending on the rules in the policy, the command might transfer just the specified Snapshot copy or Snapshot copies that are newer than the common Snapshot copy upto and including the specified Snapshot copy to the destination volume.

After the `snapmirror update` command successfully completes, the last Snapshot copy transferred is designated as the new exported Snapshot copy on the destination volume. If an update to an extended data protection (XDP) relationship specifies a Snapshot copy using the `-source-snapshot` parameter that is older than the common snapshot, after the `snapmirror update` successfully completes, the exported Snapshot copy 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 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 copy being transferred is a matching Snapshot copy or not.

If you add an aggregate to the source Infinite Volume, you must also add an aggregate of the same or greater size to the destination Infinite Volume before any `snapmirror update` occurs.

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.

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>|<hostip:/share/share-name>}] - 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 the 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-Data 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-Data ONTAP source endpoints.

{ -destination-path {<[vserver:][volume]>|<[[cluster:]//vserver/]volume>|<hostip:/lun/name>|<hostip:/share/share-name>} - 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 AltaVault destinations, the destination endpoint is specified in the form *hostip:/share/share-name*. For relationships with 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-Data 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-Data ONTAP destination endpoints.

[-s, -source-snapshot <text>] - Source Snapshot

This optional parameter specifies a Snapshot copy to transfer. The default behavior for a data protection relationship with a read-write source is that Data ONTAP creates a new Snapshot copy and uses it as the basis for determining what data are replicated; with this option, the specified Snapshot copy 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 copy must be newer than the latest common Snapshot copy. For an extended data protection relationship, the specified Snapshot copy can be newer or older than the common Snapshot copy. 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 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. 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:
```

Related Links

- snapmirror create
- snapmirror initialize
- snapmirror initialize-ls-set
- snapmirror update-ls-set
- snapmirror abort
- snapmirror show
- job show
- job history show
- snapmirror modify

snapmirror config-replication commands

snapmirror config-replication cluster-storage-configuration modify

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 (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
```

snapmirror config-replication cluster-storage-configuration show

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
```

snapmirror config-replication status show-aggregate-eligibility

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.
```

snapmirror config-replication status show-communication

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 hearbeat 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.
```

snapmirror config-replication status show

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
```

snapmirror policy commands

snapmirror policy add-rule

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 Snapshot copies are protected by vault relationships or define the schedule at which Snapshot copies are created on the SnapMirror destination. Rules which do not include a schedule are rules for protecting Snapshot copies. Rules which include a schedule are rules for creating Snapshot copies 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 copy selection for extended data protection (XDP) relationships. Only Snapshot copies 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, Snapshot copies 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 copy and transfer it as part of a 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 Snapshot copies on the primary volume along with the newly created Snapshot copy are transferred as a part of a snapmirror update or 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 Snapshot copies of the primary volume are selected for transfer.

-keep <text> - Snapshot Copy Retention Count

Specifies the maximum number of Snapshot copies that are retained on the SnapMirror destination volume for a rule. The total number of Snapshot copies retained for all the rules in a policy cannot exceed 251. 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 copy retention count is reached on the SnapMirror vault destination for the rule. The default value is *false*, which means that the oldest Snapshot copy will be deleted to make room for new ones only if the number of Snapshot copies has exceeded the retention count specified in the "keep" parameter. When set to *true*, and when the Snapshot copies have reached the retention count, then an incremental SnapMirror vault update transfer will fail or if the rule has a schedule, Snapshot copies 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 Snapshot copies (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, Snapshot copies are directly created on the SnapMirror destination. The Snapshot copies created will have the same content as the latest Snapshot copy already present on the SnapMirror destination. Snapshot copies on the source that have a SnapMirror label matching this rule will not be selected for transfer. The default value is `-`.

You define and name a schedule using the job schedule cron create command.

[-prefix <text>] - Snapshot Copy Creation Prefix

This optional parameter specifies the prefix for the Snapshot copy 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.

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* Snapshot copies.

```
vs0.example.com::> snapmirror policy add-rule -vserver vs0.example.com
       -policy TieredBackup -snapmirror-label nightly -keep 5
```

Related Links

- snapmirror update
- snapmirror initialize
- job schedule cron create

snapmirror policy create

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.

All SnapMirror policies have a field *`create-snapshot`* . This field specifies whether SnapMirror creates a new Snapshot copy on the primary volume at the beginning of a snapmirror update or 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.

Use the 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}]` - Snapmirror Policy Type

This parameter specifies the SnapMirror policy type. The supported values are *`async-mirror`* , *`vault`* and *`mirror-vault`* . Data protection (DP) relationships support only *`async-mirror`* policy type, while extended data protection (XDP) relationships support all three 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, snapmirror update and snapmirror resync operations transfer selected Snapshot copies from the primary volume to the secondary volume. The selection of Snapshot copies is governed by the rules in the policy. However snapmirror initialize and snapmirror update operations on data protection (DP) relationships ignore the rules in the policy and transfer *`all`* Snapshot copies of the primary volume which are newer than the common Snapshot copy on the destination. For both data protection (DP) and extended data protection (XDP) relationships, the Snapshot copies are kept on the secondary volume as long as they exist on the primary volume. Once a

protected Snapshot copy 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 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 Snapshot copies 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.

[-comment <text>] - Comment

Specifies a text comment for the SnapMirror policy. If the comment contains spaces, it must be enclosed within quotes.

[-tries <unsigned32_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 copy 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 copy will still be created and transferred. Data 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 copy 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* .

[-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.

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

- snapmirror update
- snapmirror resync
- snapmirror policy add-rule
- snapmirror initialize

snapmirror policy delete

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
```

snapmirror policy modify-rule

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 Snapshot copies on the vault destination when the next transfer by the snapmirror update command occurs or when the next scheduled Snapshot copy creation on the destination for the rule occurs. Modifying a rule to add a schedule will enable creation of Snapshot copies on the SnapMirror destination. Snapshot copies 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.

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 Snapshot copies that are retained on the SnapMirror destination volume for a rule. The total number of Snapshot copies retained for all the rules in a policy cannot exceed 251. 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 copy retention count is reached on the SnapMirror vault destination for the rule. The default value is *false* , which means that the oldest Snapshot copy will be deleted to make room for new ones only if the number of Snapshot copies has exceeded the retention count specified in the "keep" parameter. When set to *true* , and when the Snapshot copies have reached the retention count, then an incremental SnapMirror vault update transfer will fail or if the rule has a schedule, Snapshot copies 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 Snapshot copies (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, Snapshot copies are directly created on the SnapMirror destination. The Snapshot

copies created will have the same content as the latest Snapshot copy already present on the SnapMirror destination. Snapshot copies on the source that have a SnapMirror label matching this rule will not be selected for transfer. The default value is `-`.

> You define and name a schedule using the job schedule cron create command.

`[-prefix <text>]` - Snapshot Copy Creation Prefix

This optional parameter specifies the prefix for the Snapshot copy 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.

Examples

The following example changes the retention count for *`nightly`* Snapshot copies 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

- snapmirror update
- job schedule cron create

snapmirror policy modify

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.

> Use the 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 <unsigned32_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 copy 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 copy will still be created and transferred. Data 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 copy 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*.

[-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.

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

- snapmirror policy modify-rule

snapmirror policy remove-rule

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 Snapshot copies 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
```

snapmirror policy show

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}] - 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 <unsigned32_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.

[-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.

[-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.

[-discard-configs <network>,…] - Configurations Not Replicated During Identity Preserve Vserver DR

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 Snapshot copies 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 Snapshot copies 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 Snapshot copies.
  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.
```

snapmirror snapshot-owner commands

snapmirror snapshot-owner create

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 copy. A Snapshot copy can have at most one owner. An owner can only be added to a Snapshot copy on a read-write volume. The Snapshot copy must have a valid SnapMirror label.

Refer to the *ONTAP Data Protection Guide* for valid use cases to add an owner on a Snapshot copy.

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 copy.

[-owner <owner name>] - Snapshot Copy Owner Name

This parameter specifies the name of the owner for the Snapshot copy. 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 copy *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 copy *snap2* on volume *vol1* in Vserver *vs0.example.com*.

```
cluster1::> snapmirror snapshot-owner create -vserver vs0.example.com
         -volume vol1 -snapshot snap2
```

snapmirror snapshot-owner delete

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 copy, which was added using the 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 copy.

[-owner <owner name>] - Snapshot Copy Owner Name

This parameter specifies the name of the owner for the Snapshot copy. When not specified, the owner with

the system-generated default name will be removed.

Examples

The following example removes owner *app1* on Snapshot copy *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 copy *snap2* on volume *vol1* in Vserver *vs0.example.com*.

```
cluster1::> snapmirror snapshot-owner delete -vserver vs0.example.com
         -volume vol1 -snapshot snap2
```

Related Links

- snapmirror snapshot-owner create

snapmirror snapshot-owner show

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 Snapshot copies with owners that were added using the 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 copy.

Examples

The following example lists all Snapshot copies 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 copy *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

- snapmirror snapshot-owner create

statistics commands

statistics show-periodic

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:

This command has been deprecated and may be removed from a future version of Data 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 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 (privilege: advanced)

Selects the object for which you want to display performance data. The default object is "*`cluster`*".

[-instance <text>] - Instance (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 (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 (privilege: advanced)

If this parameter is specified, the command displays statistics for the specified preset.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the nodes for which you want to display performance data. The default node is "*`cluster:summary`*".

[-vserver <vserver name>] - Vserver (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 (privilege: advanced)

Specifies, in seconds, the interval between statistics updates. The default setting is 1 second.

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

statistics show

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 statistics start and 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] (privilege: advanced) }

If this parameter is specified, the command displays performance data in tabular format.

[-object <text>] - Object (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 statistics catalog object show . To specify multiple objects, use "|" between each object.

> 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 (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 (privilege: advanced)

Selects the counters for which you want to display performance data. To specify multiple counters, use "|" between each counter.

[-preset <text>] - Preset (privilege: advanced)

If this parameter is specified, the command displays statistics for the specified preset.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the nodes for which you want to display performance data.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Selects the Vserver for which you want to display performance data.

[-value <text>] - Text Value (privilege: advanced)

Selects the performance data that matches the specified counter value.

[-labels <text>,…] - List of Labels (privilege: advanced)

Selects the performance data that matches the specified label.

[-values <text>,…] - List of Values (privilege: advanced)

Displays only the statistics that have the specified values.

[-filter <text>] - Filter Data (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 (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 (privilege: advanced)

Specifies, in seconds, the interval between statistics updates. The default setting is 5 seconds.

[-iterations <integer>] - Iterations (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 (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 (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 (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

- statistics start
- statistics stop
- statistics catalog object show

statistics start

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 statistics stop command to stop the collection. You view the sample of performance data by using the statistics show command. You can collect more than one sample at a time.

Parameters

`[-object <text>]` - Object (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 statistics catalog object show at the command prompt. To specify multiple objects, use "|" between each object.

> 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 (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 (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 (privilege: advanced)

If this parameter is specified, the command displays statistics for the specified preset.

[-sample-id <text>] - Sample Identifier (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 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 (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 (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 (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 (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 (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 (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 (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

- statistics stop
- statistics show
- statistics catalog object show

statistics stop

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 statistics show command.

Parameters

[-sample-id <text>] - Sample Identifier (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 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

- statistics show
- statistics start

statistics aggregate commands

statistics aggregate show

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

[...]
```

statistics cache commands

statistics cache flash-pool show

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 - percent of cache 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 pools 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       Cache
                                Hit   Hit        Used    Read   Write
   Aggregate  Vserver  Volume  (%)   (%)         (%)  Blocks  Blocks
Rejects
 ----------- -------- ------- ---- ----- ----------- ------- -------
-------
       aggr1         - -total-    0     0           0       0       0
0
       aggr2      vs1    vol1    0     0           0       0       0
0
[...]
```

statistics catalog commands

statistics catalog counter show

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] (privilege: advanced) }

Displays detailed information about each counter, including privilege level, label, and whether the counter is a key counter.

-object <text> - Object (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 statistics catalog object show .

[-counter <text>] - Counter (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 (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 (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 (privilege: advanced)

Selects the counters that match this parameter value.

[-privilege <text>] - Privilegel Level (privilege: advanced)

Selects the counters that match this parameter value.

[-is-key-counter {true|false}] - Is Key Counter (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 (privilege: advanced)

Selects the counters that are deprecated (true) or are not deprecated (false).

[-replaced-by <text>] - Replaced By Counter If Deprecated (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

- statistics catalog object show

statistics catalog instance show

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 (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 statistics catalog object show .

[-instance <text>] - Instance Name (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 (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 (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 (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

- statistics catalog object show

statistics catalog object show

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] (privilege: advanced) }

Displays detailed information about each object, including privilege level.

[-object <text>] - Object (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 (privilege: advanced)

Selects the objects that match this parameter value.

[-is-deprecated {true|false}] - Is Object Deprecated (privilege: advanced)

Selects the objects that are deprecated (true) or are not deprecated (false).

[-replaced-by <text>] - Replaced By Object If Deprecated (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 (privilege: advanced)

Specifies if the object is statistically tracked

[-description <text>] - Description (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.
[...]
```

statistics disk commands

statistics disk show

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 (privilege: advanced)

Selects the disk for which you want to display performance data.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the node for which you want to display performance data.

[-sort-key <text>] - Column to Sort By (privilege: advanced)

If this parameter is specified, the command displays statistics sorted by the specified column.

-interval <integer> - Interval (privilege: advanced)

Specifies, in seconds, the interval between statistics updates. The default setting is 5 seconds.

-iterations <integer> - Iterations (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 (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

[...]
```

statistics lif commands

statistics lif show

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

[...]
```

statistics lun commands

statistics lun show

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(ms) - average latency for an operation in milliseconds.

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)     (ms)
 ---- ------- ------ ---- ----- ----- ------ ----- -------
 lun1     vs1     58   13    15    29 310585  3014      39
 lun0     vs2     56    0    11    45   8192 28826      47
[...]
```

statistics nfs commands

statistics nfs show-mount

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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 (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             -
```

statistics nfs show-nlm

(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.

This command requires an effective cluster version earlier than Data ONTAP 9.0. Data for nodes running Data ONTAP 9.0 or later is not collected, and will not be displayed. Use the 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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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

- statistics show

statistics nfs show-statusmon

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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             -
```

statistics nfs show-v3

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            -
```

statistics nfs show-v4

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>] - SetClild 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>] - SetClildConf 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 SetClild 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 SetClildConf 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             -
```

statistics node commands

statistics node show

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(ms) - average latency for an operation in milliseconds.

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    (ms)
 ----- --- ------ -------
 node2  76    113       -
 node1  58     10       -

[...]
```

statistics oncrpc commands

statistics oncrpc show-rpc-calls

Display ONC RPC Call Statistics

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

Description

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] }

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 the RPC calls performed by the node you specify.

[-protocol {TCP|UDP}] - Transport Protocol (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 (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 (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 (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 (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 (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 (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.
```

statistics port commands

statistics port fcp show

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

[...]
```

statistics preset commands

statistics preset delete

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 (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.
```

statistics preset modify

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 (privilege: advanced)

Name of the performance preset to be modified.

`[-new-name <text>]` - New Preset Name (privilege: advanced)

Set preset name to the given new name.

[-comment <text>] - Preset Description (privilege: advanced)

Set comment to the given value.

[-privilege <PrivilegeLevel>] - Preset Privilege Level (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.
```

statistics preset show

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] }

Shows details of all attributes of performance preset configuration.

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

Selects the performance presets that match the specified preset name.

[-comment <text>] - Preset Description (privilege: advanced)

Selects the performance presets that match the specified comment.

[-privilege <PrivilegeLevel>] - Preset Privilege Level (privilege: advanced)

Selects the performance presets that are available with the specified privilege.

[-is-read-only {true|false}] - Is Preset Read-Only? (privilege: advanced)

Selects the performance presets that are read-only (true) or are not read-only (false). Read-only presets cannot be modified.

[-store <text>] - Name of Store Where Data is Saved (privilege: advanced)

Selects the store where data is saved.

Examples

```
cluster1::*> statistics preset show
Preset Name         Privilege  Read-Only Comment
------------------- ---------- ---------
-------------------------------------
aggregate_overview  admin      true      This preset configuration is used
by
                                         statistics aggregate show
command.
                                         Provides overview of aggregate
object.
disk_overview       advanced   true      This preset configuration is used
by
                                         statistics disk show command.
                                         Provides overview of disk object.
fcp_port_overview   admin      true      This preset configuration is used
by
                                         statistics port fcp show command.
                                         Provides overview of fcp port
object.
flash_pool_overview admin      true      This preset configuration is used
by
                                         statistics cache flash-pool show
                                         command. Provides overview of
flash
                                         pool object.
[...]
```

statistics preset detail show

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] }

Displays all of the performance preset detail attributes.

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

Selects the performance preset details that match the specified preset name.

[-object <text>] - Performance Object (privilege: advanced)

Selects the performance preset details that match the specified object name.

[-sample-period <sample_period>] - Archive Sample period (privilege: advanced)

Selects the performance preset details that are collected at the specified sample period.

[-counter-set <text>,…] - Performance Counter Name Set (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 (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,
. . .
```

statistics samples commands

statistics samples delete

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 statistics start command.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Selects the Vserver for which you want to delete the sample. The default Vserver is admin Vserver.

`-sample-id <text>` - Sample Identifier (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

- statistics start

statistics samples show

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 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 ]` (privilege: advanced) }

Displays detailed information about each sample.

[-vserver <vserver name>] - Vserver (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 (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

- statistics start

statistics settings commands

statistics settings modify

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 (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 (privilege: advanced)

Specifies whether `statistics` commands display per-client information. The default is disabled.

If you enable this setting, you might significantly impact system performance.

`[-counter-display-units {B|KB|MB|GB}]` - Counter Display Units (privilege: advanced)

Specifies display units for the counters. The default setting is MB.

`[-display-count-exponent <integer>]` - Display Count Exponent (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
```

statistics settings show

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
```

statistics system commands

statistics system show

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       -

[...]
```

statistics top commands

statistics top client show

Most active clients

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

Description

This command continuously displays performance data for top clients at a regular interval. The command output displays data in the following columns:

- Client - client name.
- Vserver - vserver name.
- Node - node name.
- Protocol - protocol name.
- Total Ops - 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

[...]
```

statistics top file show

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:

- File - file name.
- Volume - volume name.
- Vserver - vserver name.
- Aggregate - aggregate name.
- Node - node name.
- Total Ops - 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 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

                                                                    *Total
Total
                  File  Volume    Vserver Aggregate         Node     Ops
(Bps)
--------------------- -------  --------- --------- ------------- ------
-----
/vol/vol01/clus/cache   vol01  vserver01     aggr1 cluster-node2      9
80
           /vol/vol02   vol02  vserver02     aggr2 cluster-node1      6
50

[...]
```

statistics volume commands

statistics volume show

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
[...]
```

statistics vserver commands

statistics vserver show

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

[...]
```

statistics workload commands

statistics workload show

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 (privilege: advanced)

Selects the workload for which you want to display performance data.

`[-sort-key <text>]` - Column to Sort By (privilege: advanced)

If this parameter is specified, the command displays statistics sorted by the specified column.

`-interval <integer>` - Interval (privilege: advanced)

Specifies, in seconds, the interval between statistics updates. The default setting is 5 seconds.

`-iterations <integer>` - Iterations (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 (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       -

[...]
```

statistics-v1 commands

statistics-v1 nfs commands

statistics-v1 nfs show-mount

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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 (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             -
```

statistics-v1 nfs show-nlm

(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.

> This command requires an effective cluster version earlier than Data ONTAP 9.0. Data for nodes running Data ONTAP 9.0 or later is not collected, and will not be displayed. Use the 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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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

- statistics show

statistics-v1 nfs show-statusmon

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

If you specify this parameter, the command displays statistics only for the specified node.

[-result {success|failure|all}] - Result (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 (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 (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 (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 (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 (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 (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 (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 (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             -
```

statistics-v1 nfs show-v3

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            -
```

statistics-v1 nfs show-v4

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>] - SetClId 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>] - SetClIdConf 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 SetClild 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 SetClildConf 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              -
```

statistics-v1 protocol-request-size commands

statistics-v1 protocol-request-size show

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            -
```

storage-service commands

storage-service show

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.

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

Selects the available Storage Services for Vservers that match the parameter value.

[-storage-service <text>] - Storage Service (privilege: advanced)

Selects the available Storage Services whose name matches the parameter value.

[-description <text>] - Description (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 (privilege: advanced)

Selects the available Storage Services whose expected IOPs per TB matches the parameter value. This field is the number of IOPs per provisioned TB nominally guaranteed by this Storage Service.

[-peak-iops-per-tb <integer>] - Peak IOPS per TB (privilege: advanced)

Selects the available Storage Services whose peak IOPs per TB matches the parameter value. This field is the number of IOPs per provisioned TB used as the maximum Quality of Service (QoS) throttle.

[-absolute-min-iops <integer>] - Absolute Minimum IOPS (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) (privilege: advanced)

Selects the available Storage Service whose target latency matches the parameter value.

[-aggr-list <aggregate name>,…] - Aggregate List (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.

storage commands

storage aggregate commands

storage aggregate add-disks

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}]` - 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.

Only the aggregates marked as `hybrid-enabled` can be converted to Flash Pools. Use storage aggregate modify command to mark the aggregate as `hybrid-enabled` .

When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity` or `raid.mix.hdd.disktype.performance` . 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.
- array = Logical storage devices backed by storage arrays and used by Data ONTAP as disks. Includes

disk type LUN.

- virtual = Virtual disks that are formatted and managed by the hypervisor. Includes disk type VMDISK.

> When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity` or `raid.mix.hdd.disktype.performance` .

[-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, Data 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 Data 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 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 Data 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.

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 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 raid4
```

Related Links

- storage aggregate modify
- storage pool show-available-capacity

storage aggregate auto-provision

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 the aggregates should be created as recommended. On positive response, ONTAP creates 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.

The command is not affected by the CLI session setting: 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:
                                                                  -Devices To
Use-
Node                New Data Aggregate            Usable Size Disks
Partitions
------------------ ---------------------------- ------------ -----
----------
node1               node1_SSD_1                        3.66TB     5
48
node2               node2_SSD_1                        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   Remaining
Node               Type    Usable Size Partition    Spares
------------------ ------ ------------ --------- ---------
node1              SSD         185.5GB disk              1
node2              SSD         185.5GB disk              1

Do you want to create recommended aggregates? {y|n}: y

Info: Creating node1_SSD_1 ...
      Creating node2_SSD_1 ...
```

Related Links

- set

storage aggregate create

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, the maximum number of disks that can be included in a RAID group, and whether the aggregate's contents are encrypted.

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}] - Disk Type

This parameter specifies the type of disk on which the aggregate is to be created.

> When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity` or `raid.mix.hdd.disktype.performance` . 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.
- array = Logical storage devices backed by storage arrays and used by Data ONTAP as disks. Includes disk type LUN.
- virtual = Virtual disks that are formatted and managed by the hypervisor. Includes disk type VMDISK.

> When this parameter is used, disk selection is not influenced by RAID options `raid.mix.hdd.disktype.capacity` or `raid.mix.hdd.disktype.performance` .

[-m, -mirror <true>] - Mirror

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, Data 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 Data 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, Data ONTAP selects the node where the aggregate is created.

[-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 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.

[-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 (privilege: advanced)

This parameter can be used to force the creation of a 2-disk RAID4 aggregate, or a 3-disk or 4-disk RAID-DP aggregate.

[-is-autobalance-eligible {true|false}] - Is Eligible for Auto Balance Aggregate (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* .

[-encrypt <true>] - Enable Encryption

This parameter specifies that the new aggregate be encrypted. If this parameter is set to *true* , the specified aggregate's contents will be encrypted.

[-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 (%) (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 (%) (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.

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 -nodes 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 -nodes 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 -nodes 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
```

The following example creates an aggregate named aggr1 on the local node. The aggregate contains 3 disks and is encrypted

```
cluster1::> storage aggregate create -aggregate aggr1 -diskcount 3
-encrypt true
```

storage aggregate delete

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 configuraton 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 (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
```

storage aggregate mirror

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
```

storage aggregate modify

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}] - Disk Type

This parameter specifies either the HDD tier or the SSD tier when changing the RAID type of a Flash Pool. 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.

[-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 Data 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 Snapshot copies 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 FlexVols that cannot be write cached. A FlexVol cannot be write-cached if it is part of an aggregate created in Data 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 FlexVols that cannot be write cached. A FlexVol cannot be write-cached if it is part of an aggregate created in Data ONTAP 7. Use `-force-hybrid-enabled` to over-ride this behavior. Note that read caching will be enabled on these FlexVols, 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.

For Flash Pools, 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.

This parameter is applicable only for Flash Pools.

[-t, -raidtype {raid_tec|raid_dp|raid4}] - 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 60 (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 (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 (%) (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 (%) (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 Data 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 perfomance by optimizing the amount of data nvlogged by user writes on platforms where NVRAM and secondary storage are of same media type.

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
```

storage aggregate offline

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
```

storage aggregate online

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
```

storage aggregate remove-stale-record

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 (privilege: advanced)

This parameter specifies the aggregate that corresponds to the stale aggregate record that is to be deleted.

-nodename {<nodename>|local} - Node Name (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
```

storage aggregate rename

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
```

storage aggregate restrict

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
```

storage aggregate scrub

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.

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%
```

storage aggregate show-cumulated-efficiency

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.

| [-details]

Use this parameter to show additional Storage Efficiency Ratios.

| [-all-details] (privilege: advanced)

Use this parameter to show additional Storage Efficiency Ratios and size values.

| [-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 (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 (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 {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used (privilege: advanced)

Displays the total logical size used in all the specified aggregates excluding Snapshots and FlexClones.

[-total-data-reduction-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used (privilege: advanced)

Displays the total physical size used by all the specified aggregates aggregates excluding Snapshots and FlexClones

[-total-data-reduction-efficiency-ratio <text>] - Total Data Reduction Efficiency Ratio

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction and Pattern Detection data reduction technologies on the specified aggregates.

[-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Space Used for All volumes (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 (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 (privilege: advanced)

Displays the total disk space that is saved by deduplication, Zero pattern detection and FlexClone for files or LUNs 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 or LUNs by all volumes in all the specified aggregates.

[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by volume Compression (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 (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 (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 (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 (privilege: advanced)

Displays the physical size used by all Volume Snapshots residing in all the specified aggregates.

[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by FlexClone volumes (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 (privilege: advanced)

Displays the physical size used by all FlexClone volumes in all the specified aggregates.

[-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.

[-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-data-reduction-storage-efficiency-ratio <text>] - FlexClone volume Data Reduction Ratio

Displays the 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 (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 Storage Efficiency Ratio: 6.97:1

cluster::> aggr show-cumulated-efficiency -details
               Total Storage Efficiency Ratio: 6.97:1
                   Total Data Reduction Ratio: 8.44:1
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 Storage Efficiency Ratio: 7.41:1
saiscluster-1::*> aggr show-cumulated-efficiency -all-details
----- Total Storage Efficiency ----------------
    Logical    Physical                 Storage
       Used        Used        Efficiency Ratio
----------- ----------- ----------------------
    89.11MB     12.91MB                  6.90:1

----- Total Data Reduction Efficiency ---------
    Logical    Physical                 Storage
       Used        Used        Efficiency Ratio
----------- ----------- ----------------------
```

```
    89.11MB     10.69MB                  8.34:1

-- 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 Deduplicaiton Disabled Volumes: 2
```

storage aggregate show-efficiency

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.

| [-details]

Use this parameter to show additional Storage Efficiency Ratios.

| [-advanced] (privilege: advanced)

Use this parameter to show additional Storage Efficiency Ratios and size values.

| [-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 (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 (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 {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Logical Used (privilege: advanced)

Displays the logical size used in the aggregate excluding Snapshots and FlexClones.

[-total-data-reduction-physical-used {<integer>[KB|MB|GB|TB|PB]}] - Total Data Reduction Physical Used (privilege: advanced)

Displays the physical size used by the aggregate excluding Snapshots and FlexClones

[-total-data-reduction-efficiency-ratio <text>] - Total Data Reduction Efficiency Ratio

Displays the total storage efficiency ratio obtained by Deduplication, Compression, Data Compaction and Pattern Detection data reduction technologies 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 (privilege: advanced)

Displays the total disk space that is saved by deduplication and FlexClone for files or LUNs 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 (privilege: advanced)

Displays the total disk space that is saved by deduplication, Zero pattern detection and FlexClone for files or LUNs 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 or LUNs by all volumes in the aggregate.

[-volume-compression-saved {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by Volume Compression (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 (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 (privilege: advanced)

Displays the physical size used by the aggregate.

[-aggr-compaction-saved {<integer>[KB|MB|GB|TB|PB]}] - Space Saved by Aggregate Data Reduction (privilege: advanced)

Displays the total disk space that is saved by data compaction at the aggregate level.

[-aggr-data-reduction-storage-efficiency-ratio <text>] - Aggregate Data Reduction SE Ratio (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 (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 (privilege: advanced)

Displays the physical size used by all Volume Snapshots residing in the aggregate.

[-flexclone-volume-logical-used {<integer>[KB|MB|GB|TB|PB]}] - Logical Size Used by FlexClone Volumes (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 (privilege: advanced)

Displays the physical size used by all FlexClone volumes in 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.

[-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-data-reduction-storage-efficiency-ratio <text>] - FlexClone Volume Data Reduction Ratio

Displays the 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 (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 Storage Efficiency Ratio: 4.29:1
Aggregate: aggr2
     Node: node1

Total Storage Efficiency Ratio: 4.29:1
cluster::*> aggr show-efficiency -details

Aggregate: aggr1
     Node: node1

Total Storage Efficiency Ratio:                4.29:1
Total Data Reduction Ratio:                    2.39:1

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:           -
Number of Efficiency Disabled Volumes:         1

Aggregate: aggr2
     Node: node1
Total Storage Efficiency Ratio:                4.29:1
Total Data Reduction Ratio:                    2.39:1

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:           -
Number of Efficiency Disabled Volumes:         1
```

storage aggregate show-resync-status

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.
```

storage aggregate show-scrub-status

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
```

storage aggregate show-space

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 Snapshot copies. This field includes the space that is reserved for Snapshot copies 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 Snapshot copies.

[-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 accomodate 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 Snapshot copies.

[-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.

[-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.

[-object-store-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 whose physical space use in the attached object store matches the specified value.

[-object-store-physical-used-percent <percent_no_limit>] - 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 Referenced Capacity

If this parameter is specified, the command displays information only about the aggregates whose reference capacity space in use in the attached object store matches the specified value.

[-object-store-referenced-capacity-percent <percent_no_limit>] - Object Store Referenced Capacity Percentage

If this parameter is specified, the command displays information only about aggregates whose 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]}] - Object Store Metadata

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>] - Object Store Metadata Percent

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]}] - Total Unreclaimed Space

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>] - Object Store Unreclaimed Space Percentage

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.

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%
Total Physical Used                     811.2MB         0%
```

storage aggregate show-spare-disks

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] (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}] - 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}] - 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.

[-standard-disk-type {SATA | FC | NL-SAS | LUN | SAS | SCSI | SSD | VM-DISK | NVMe-SSD}] - 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.
- array — Logical storage devices backed by storage arrays and used by Data 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 Data 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 (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 (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 (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 (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.
```

storage aggregate show-status

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.

> This command does not use pagination. You can reduce the output by filtering with the parameters 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] }

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)
```

storage aggregate show

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>] - Mirror

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 Snapshot copies.

[-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 (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.

For Flash Pools, 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 Pools.

This parameter is applicable only for Flash Pools.

[-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 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 Data 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 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 (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 (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 (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 Snapshot copies.

[-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 (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.

[-is-encrypted {true|false}] - Is Encrypted

Selects the aggregates that are encrypted.

[-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.

[-encryption-key-id <text>] - Encryption Key ID

Selects the aggregates that are encrypted with the specified key ID.

[-is-cft-precommit {true|false}] - Is in the precommit phase of Copy-Free Transition (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 (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 clustered Data ONTAP, and are currently out of

space.

[-autobalance-unbalanced-threshold-percent <integer>] - Threshold When Aggregate Is Considered Unbalanced (%) (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 (%) (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 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 Data 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 FabricPool matches the specified value. A FabricPool has an 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 information only about aggregates whose space in use in the attached external capacity tier matches the specified value.

[-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.

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
Data 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

- storage aggregate show-space
- storage aggregate resynchronization modify

storage aggregate verify

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.

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
```

storage aggregate efficiency show

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 has 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 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.
```

storage aggregate efficiency cross-volume-dedupe revert-to

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 (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 (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" successfully.

cluster:::> storage aggregate efficiency cross-volume-dedupe revert-to
-aggregate aggr1 -clean-up true
The revert operation started on aggregate "aggr1" successfully.
```

storage aggregate efficiency cross-volume-dedupe show

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.
```

storage aggregate efficiency cross-volume-dedupe stop

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 (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.
```

storage aggregate inode-upgrade resume

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 (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 (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
```

storage aggregate inode-upgrade show

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] }

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

[-aggregate <text>] - Aggregate Name (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 (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 (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 (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 (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>h][<integer>m][<integer>s]>] - Remaining Upgrade Time (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 (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
```

storage aggregate object-store attach

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.

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.

Examples

The following example attaches an object store to aggregate aggr1:

```
cluster1::>storage aggregate object-store attach -aggregate aggr1 -object
-store-name my-store
```

storage aggregate object-store show-freeing-status

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 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] }

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

[-bin-uuid <UUID>] - UUID of the Bin (privilege: advanced)

If this parameter is specified, the command displays information only about the aggregate attached to the specified bin UUID.

[-object-store-name <text>] - Object Store Configuration Name (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 (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}] - Request State (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 (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 (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

- storage aggregate delete

storage aggregate object-store show-space

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.

[-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 specfied value. If the object store does not require a license, then this field is not set.

Examples

The following example displays space information about all object stores:

```
cluster1::>storage aggregate object-store show-space
```

storage aggregate object-store show

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.

[-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.

Examples

The following example displays all information about all object stores:

```
cluster1::>storage aggregate object-store show
```

storage aggregate object-store config create

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 Data ONTAP how to connect to an object store. Following pre-requisites must be met before creating an object store configuration in Data 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 Data 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. Data 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 object store).

An object-store configuration once created must not be reassociated with a different object-store or container. See storage aggregate object-store config modify command for more information. If no secret password or access key is provided while setting up a configuration for AWS_S3 object store, Data ONTAP will try to obtain them through Data ONTAP Cloud in AWS using the AWS Identity and Access Management (IAM) role. If Data 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) and SGWS (StorageGrid WebScale).

`-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'.

[-ssl-enabled {true|false}] - SSL/HTTPS Enabled?

This parameter indicates whether secured connection will be used during data access to the object store. Default: False. Use of SSL certificates is one such method to ensure secure communication.

[-port <integer>] - Port Number of the Object Store

This parameter specifies the port number on the remote server that Data ONTAP will use while establishing connection to the object store.

-container-name <text> - Data Bucket/Container Name Used in S3 Requests

This parameter specifies the data bucket or container that Data ONTAP should read and write to.

[-access-key <text>] - Access Key

This parameter specifies the access key (access key ID) required to authorize requests to the object store.

[-secret-password <text>] - Secret Password

This parameter specifies the password (secret access key) to authenticate requests to the 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*

[-use-iam-role {true|false}] - Use IAM Role for AWS ONTAP Cloud

This parameter specifies whether IAM credentials must be used for data access from AWS_S3 object store. This parameter is required only when setting up a AWS_S3 object store configuration in Data ONTAP. The IAM credentials required for setting up this configuration will obtained through Data ONTAP Cloud in AWS. This option is available only through ONTAPI.

Examples

The following example creates a object store configuration in Data ONTAP:

```
cluster1::>storage aggregate object-store config create -object-store-name
my-aws_store
           -provider-type AWS_S3 -server objstr_srvr.amazon.com -port 1237
-s3-name AWS-bucket
           -ssl-enabled false -ipspace Default
```

Related Links

- storage aggregate object-store config modify

storage aggregate object-store config delete

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 Data 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 storage aggregate delete command. The command storage aggregate object-store show can be used to view the aggregates attached to the object store before issuing the delete command.

The 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

- storage aggregate delete
- storage aggregate object-store show

storage aggregate object-store config modify

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 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.

[-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'.

[-ssl-enabled {true|false}] - SSL/HTTPS Enabled?

This optional parameter can be toggled to indicate whether secured connection should be used for communication with the object store server. Use of SSL certificates is one such method to ensure secure communication.

[-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

This optional parameter specifies a new access key for the object store.

[-secret-password <text>] - Secret Password

This optional parameter specifies a new password for the object store.

[-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}] - Use IAM Role for AWS ONTAP Cloud

This optional parameter specifies whether to enable IAM role for this object store configuration. IAM role can be enabled only if connecting to AWS object store in Data ONTAP Cloud. When using this parameter, `-secret-password` and `-access-key` parameters must not be explicitly specified as they will be obtained from Data ONTAP Cloud in AWS using the IAM role.

Examples

The following example modifies two parameters (port number and ssl-enabled) of an object store configuration named my-store:

```
cluster1::>storage aggregate object-store config modify -object-store-name
my-store -port 1235 -ssl-enabled true
```

storage aggregate object-store config rename

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 paramter 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
```

storage aggregate object-store config show

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.

[-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).

[-ssl-enabled {true|false}] - SSL/HTTPS 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 (for example, SSL certificates) 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 Used in S3 Requests

If this parameter is specified, the command displays information only about object store configurations whose container name matches the specified value. Data ONTAP uses this container name or object store data bucket while accessing data from the object store.

[-access-key <text>] - Access Key

If this parameter is specified, the command displays information only about object store configurations whose access key matches the specified value. Data 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. Data ONTAP uses the IPspace value to connect to the object store.

[-use-iam-role {true|false}] - Use IAM Role for AWS ONTAP Cloud

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 Data ONTAP Cloud.

[-iam-role <text>] - IAM Role for AWS ONTAP Cloud

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.

Examples

The following example displays all available object store configuration in the cluster:

```
cluster1::>storage aggregate object-store config show
```

storage aggregate plex delete

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
```

storage aggregate plex offline

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
```

storage aggregate plex online

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
```

storage aggregate plex show

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
```

storage aggregate reallocation quiesce

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 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

- storage aggregate reallocation restart

storage aggregate reallocation restart

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.

storage aggregate reallocation schedule

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 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 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.

> 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

- storage aggregate reallocation start

storage aggregate reallocation show

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|Reschedule|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 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

- job schedule show

storage aggregate reallocation start

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 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.

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

- storage aggregate reallocation schedule

storage aggregate reallocation stop

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.

storage aggregate relocation show

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.
```

storage aggregate relocation start

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 Data 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 Data 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 (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
```

storage aggregate resynchronization modify

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 Data 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
```

storage aggregate resynchronization show

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 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 Data 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

- storage aggregate resynchronization options show

storage aggregate resynchronization options modify

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 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

- storage aggregate resynchronization modify

storage aggregate resynchronization options show

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
```

storage array commands

storage array modify

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 Data 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 (privilege: advanced)

The target port queue depth for all target ports on this array.

[-lun-queue-depth <integer>] - LUN Queue Depth (privilege: advanced)

The queue depth assigned to array LUNs from this array.

{ [-is-upgrade-pending {true|false}] - Upgrade Pending (privilege: advanced)

Set this parameter to *true* if the array requires additional Data 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 Data ONTAP declares an "all path failure" event for an array. Delaying the "all path failure" event allows Data 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
```

storage array remove

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. Upon command completion, if a storage array is still connected to the cluster, the array profile record is re-created with default values.

Parameters

`-name <text>` - Name

Name of the storage array you want to remove from the database.

Examples

```
cluster1::> storage array remove IBM_1722_1
```

storage array rename

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 Data ONTAP assigned during device discovery. By default, the name that Data ONTAP assigned to the storage array during discovery is shown in Data ONTAP displays and command output.

Parameters

`-name <text>` - Name

Storage array name either generated by Data 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
```

storage array show

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 (privilege: advanced)

Selects the arrays that match this parameter value.

[-lun-queue-depth <integer>] - LUN Queue Depth (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 (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
```

storage array config show

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.
```

storage array disk paths show

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: `20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000
`

[-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 (Assymetric 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:00000000:00000000:00000000:00000000:00000000:00000000:00
000000:00000000
                                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:
5000C500:0979E3C3:00000000:00000000:00000000:00000000:00000000:00000000:00
000000:00000000
                                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:
5000C500:0979E3F3:00000000:00000000:00000000:00000000:00000000:00000000:00
000000:00000000
                                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:00000000:00000000:00000000:00000000:00000000:00000000:00
000000:00000000
                                  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
[...]
```

storage array port modify

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 Data 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
```

storage array port remove

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 Data 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
```

storage array port show

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 Data 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 (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 (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 (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 (privilege: advanced)

Current dynamic target port queue depth, the maximum number of commands allowed outstanding.

[-max-pending <integer>,…] - Max Pending (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 (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
```

storage automated-working-set-analyzer commands

storage automated-working-set-analyzer show

Display running instances

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

Description

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

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[-cacheable-write <percent>] - Cacheable Write (privilege: advanced)

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

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

This parameter indicates the projected Flash Pool cache usage.

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

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

[-projected-write-hit <percent>] - Projected Write Hit (privilege: advanced)

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

[-referenced-interval-id <integer>] - Referenced Interval ID (privilege: advanced)

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

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

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

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

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

[-read-hit-20 <percent>] - 20% Cache Read Hit (privilege: advanced)

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

[-read-hit-40 <percent>] - 40% Cache Read Hit (privilege: advanced)

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

[-read-hit-60 <percent>] - 60% Cache Read Hit (privilege: advanced)

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

[-read-hit-80 <percent>] - 80% Cache Read Hit (privilege: advanced)

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

[-read-hit-100 <percent>] - 100% Cache Read Hit (privilege: advanced)

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

size.

[-write-hit-20 <percent>] - 20% Cache Write Hit (privilege: advanced)

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

[-write-hit-40 <percent>] - 40% Cache Write Hit (privilege: advanced)

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

[-write-hit-60 <percent>] - 60% Cache Write Hit (privilege: advanced)

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

[-write-hit-80 <percent>] - 80% Cache Write Hit (privilege: advanced)

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

[-write-hit-100 <percent>] - 100% Cache Write Hit (privilege: advanced)

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

Examples

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

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

storage automated-working-set-analyzer start

Command to start Automated Working Set Analyzer on node or aggregate

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

Description

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

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

- Workload monitoring

Parameters

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

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

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

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

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

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

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

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

Examples

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

storage automated-working-set-analyzer stop

Command to stop Automated Working Set Analyzer on node or aggregate

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

Description

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

Parameters

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

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

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

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

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

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

Examples

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

storage automated-working-set-analyzer volume show

Displays the Automated Working Set Analyzer volume table

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

Description

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

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[-cacheable-write <percent>] - Cacheable Write (privilege: advanced)

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

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

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

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

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

[-projected-write-hit <percent>] - Projected Write Hit (privilege: advanced)

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

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
```

storage bridge commands

storage bridge add

Add a bridge for monitoring

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

Description

The `storage bridge add` command enables you to add FC-to-SAS bridges for SNMP monitoring in a MetroCluster configuration.

Parameters

-address <IP Address> - Bridge Management Port IP Address

This parameter specifies the IP address of the bridge that is being added for monitoring.

[-snmp-community <text>] - SNMP Community

This parameter specifies the SNMP community set on the bridge that is being added for monitoring.

[-veto-backend-fabric-check {true|false}] - Veto Backend Fabric Check (privilege: advanced)

If specified, the `storage bridge add` command will not check if the bridge is present in the MetroCluster's backend fabric. By default, it does not let you add bridges that are not present.

Examples

The following command adds a bridge with IP address '10.226.197.16' for monitoring:

```
cluster1::> storage bridge add -address 10.226.197.16
cluster1::> storage bridge show
                                                                    Is
Monitor
    Bridge     Symbolic Name Vendor  Model     Bridge WWN       Monitored
Status
    ---------- ------------- ------- --------- ---------------- ---------
-------
    ATTO_10.226.197.16
               Bridge Number 16
                             Atto    FibreBridge 6500N
                                                 2000001086603824 true
-
    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.
cluster1::>
```

storage bridge modify

Modify a bridge's configuration information

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

Description

The `storage 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.

`[-address <IP Address>]` - Bridge IP Address

This parameter specifies the IP address of the bridge.

`[-snmp-community <text>]` - SNMP Community Set on the Bridge

This parameter specifies the SNMP community set on the bridge.

Examples

The following command modifies 'ATTO_10.226.197.16' bridge SNMP community to 'public':

```
cluster1::> storage bridge modify -name ATTO_10.226.197.16 -address
10.226.197.16 -snmp-community public
cluster1::>
```

storage bridge refresh

Refresh storage bridge info

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

Description

The `storage bridge refresh` command triggers a refresh of the SNMP data for the MetroCluster 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 storage switch add and storage bridge add commands respectectively.

Examples

The following command triggers a refresh for the SNMP data:

```
cluster1::*> storage bridge refresh
cluster1::*>
```

Related Links

- storage switch add
- storage bridge add

storage bridge remove

Remove a bridge from monitoring

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

Description

The `storage 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::> storage bridge remove -name ATTO_10.226.197.16
cluster1::> storage 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.
```

storage bridge show

Display bridge information

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

Description

The `storage 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 storage bridge add command. If no parameters are specified, the default command displays the following information about the storage bridges:

- Bridge
- Symbolic Name
- Vendor
- Model
- Bridge WWN

- Is Monitored
- Monitor Status

To display detailed profile information about a single storage bridge, use the `-name` parameter.

Parameters

{ [-fields <fieldname>,…]

Displays the specified fields for all the storage bridges, in column style output.

| [-connectivity]

Displays the following details about the connectivity from different entities to the storage 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 storage bridge:

- Sensor Name
- Reading in degree Celsius ©
- Fan operational status
- Minimum Safe Operating Temperature in degree Celsius ©
- Maximum Safe Operating Temperature in degree Celsius ©
- Sensor Status

| [-error]

Displays the errors related to the storage bridge.

| [-ports]

Displays the following details about the storage 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 storage 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 7500 only:

- Power supply name
- Power supply status

| [-sfp]

Displays the following details about the storage 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 storage 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 storage 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 storage 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 storage 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 storage bridges in the system. If a storage bridge is specified, then this parameter displays the same detailed information for the storage bridge you specify as does the -name parameter.

[-name <text>] - Bridge Name

Displays information only about the storage bridges that match the name you specify.

[-wwn <text>] - Bridge World Wide Name

Displays information only about the storage bridges that match the bridge wwn you specify.

[-model <text>] - Bridge Model

Displays information only about the storage bridges that match the bridge model you specify.

[-vendor {unknown|Atto}] - Bridge Vendor

Displays information only about the storage bridges that match the bridge vendor you specify.

[-fw-version <text>] - Bridge Firmware Version

Displays information only about the storage bridges that match the bridge firmware version you specify.

[-serial-number <text>] - Bridge Serial Number

Displays information only about the storage bridges that match the bridge serial number you specify.

[-address <IP Address>] - Bridge IP Address

Displays information only about the storage bridges that match the bridge IP address you specify.

[-is-monitoring-enabled {true|false}] - Is Monitoring Enabled for Bridge?

Displays information only about the storage bridges that match the bridge monitoring value you specify.

[-status {unknown|ok|error}] - Bridge Status

Displays information only about the storage 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 storage bridges that match the profile data last successful refresh timestamp you specify.

[-symbolic-name <text>] - Bridge Symbolic Name

Displays information only about the storage bridges that match the symbolic name you specify.

[-snmp-community <text>] - SNMP Community Set on the Bridge

Displays information only about the storage bridges that match the bridge SNMP community you specify.

[-error-text-list <text>,…] - Bridge Error Description List

Displays information only about the storage bridges that have the errors you specify.

[-temp-sensor-name <text>] - Temperature Sensor Name

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage bridges that have the FC ports with the SFP part numbers you specify.

[-fc-port-sfp-data-rate-capability-list {2Gb|4Gb|8Gb|16Gb}] - Bridge FC Port SFP Data Rate Capability List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage bridges that are connected to the nodes you specify.

[-initiator-list <text>,…] - Initiator List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage bridges that match the target ports with indexes you specify.

Examples

The following example displays information about all storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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 storage bridges:

```
cluster1::> storage 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

- storage bridge add

storage bridge config-dump collect

Retrieve and save bridge dumpconfiguration

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

Description

The `storage bridge config-dump collect` command retrieves a dumpconfiguration file from a storage bridge.

Parameters

`-bridge <text>` - Bridge Name (privilege: advanced)

Use this parameter to retrieve a dumpconfiguration file from the specified storage bridge.

Examples

The following example retrieves a dumpconfiguration file from storage bridge ATTO_FibreBridge7500N_1:

```
cluster1::*> storage bridge config-dump collect -bridge
ATTO_FibreBridge7500N_1
[Job 883] Job is queued: Collect the dumpconfiguration file from bridge
"ATTO_FibreBridge7500N_1".

cluster1::*>
```

storage bridge config-dump delete

Delete a dumpconfiguration file

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

Description

The `storage bridge config-dump delete` command deletes dumpconfiguration files previously retrieved with the storage bridge config-dump collect command.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to delete a dumpconfiguration file stored on the specified node.

`-file <text>` - Config File (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::*> storage bridge config-dump delete -node node1 -file
dsbridge_config.FB7500N100001.2017-04-28_14_49_30.txt

cluster1::*>
```

Related Links

- storage bridge config-dump collect

storage bridge config-dump show

Display a list of bridge dumpconfiguration files

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

Description

The `storage bridge config-dump show` command displays information about all the dumpconfiguration files previously retrieved with the storage 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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Displays information about the dumpconfiguration files stored on the node that matches the specified node name.

[-file <text>] - Config File (privilege: advanced)

Displays information about the dumpconfiguration files that match the specified file name.

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

Displays information about the dumpconfiguration files from the storage bridge that matches the specified bridge name.

[-serial-number <text>] - Serial Number of Bridge (privilege: advanced)

Displays information about the dumpconfiguration files from the storage bridge that matches the specified serial number.

[-timestamp <MM/DD/YYYY HH:MM:SS>] - Time of Collection (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::*> storage 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::*> storage 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

- storage bridge config-dump collect

storage bridge coredump collect

Retrieve and save coredump

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

Description

The `storage bridge coredump collect` command retrieves a core file from a storage bridge.

Parameters

`-name <text>` - Bridge Name

This parameter specifies the storage bridge name from which the coredump file is to be collected.

Examples

The following example retrieves a coredump from storage bridge ATTO_FibreBridge7500N_1:

```
cluster1::> storage bridge coredump collect -bridge
ATTO_FibreBridge7500N_1
[Job 883] Job is queued: Collect the coredump from bridge
"ATTO_FibreBridge7500N_1".

cluster1::>
```

storage bridge coredump delete

Delete a saved coredump file.

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

Description

The `storage bridge coredump delete` command deletes a coredump file previously retrieved with the storage 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::> storage bridge coredump delete -name ATTO_FibreBridge7500N_1
-corename core.FB7500N100018.1970-01-05.17_50_30.mem

cluster1::>
```

Related Links

- storage bridge coredump collect

storage bridge coredump show

Display a list of bridge coredumps

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

Description

The `storage bridge coredump show` command displays information about all the coredump files previously retrieved with the storage 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 storage 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 storage 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::> storage 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

- storage bridge coredump collect

storage disk commands

storage disk assign

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 Data 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 Data 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.

| -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}] - 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, regardless of the setting of the disk.auto_assign option. 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 Pool0 or Pool1.

{ [-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 Data 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 (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 (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 (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 (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
```

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
```

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
```

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
```

storage disk fail

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 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

- storage disk unfail

storage disk reassign

(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 Data 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 (privilege: advanced)

This specifies the serial number of the failed node.

`-d, -newhomeid <nvramid>` - New Home ID (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 (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.
```

storage disk refresh-ownership

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 (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
```

storage disk remove-reservation

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 (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
```

storage disk remove

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 storage disk removeowner command instead. See the "Physical Storage Management Guide" for the complete procedure.

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 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

- storage disk removeowner
- storage disk replace

storage disk removeowner

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.

{ `[-root <true>]` - Root Partition of Root-Data/Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter removes ownership of 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 (privilege: advanced)

This optional parameter removes ownership of the data partition of a root-data partitioned disk. You cannot use this parameter with a root-data1-data2 partitioned disk or disks that are part of a storage pool. The default value is *`false`* .

| `[-data1 <true>]` - Data1 Partition of a Root-Data1-Data2 Partitioned Disk (privilege: advanced)

This optional parameter removes ownership of the data1 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with a root-data partitioned disk or disks that are part of a storage pool. The default value is *`false`* .

| `[-data2 <true>]` - Data2 Partition of a Root-Data1-Data2 Partitioned Disk (privilege: advanced) }

This optional parameter removes ownership of the data2 partition of a root-data1-data2 partitioned disk. You cannot use this parameter with a root-data partitioned disk or disks that are part of a storage pool. The default value is *`false`* .

Examples

The following example removes the ownership from a disk named 1.1.27.

```
cluster1::> storage disk removeowner -disk 1.1.27
```

The following example removes ownership of the root partition on disk 1.1.16.

```
cluster1::> storage disk removeowner -disk 1.1.16 -root true
```

The following example removes ownership of the data partition on disk 1.1.16.

```
cluster1::> storage disk removeowner -disk 1.1.16 -data true
```

The following example removes ownership of the data1 partition on disk 1.1.23.

```
cluster1::> storage disk removeowner -disk 1.1.23 -data1 true
```

The following example removes ownership of the data2 partition on disk 1.1.23.

```
cluster1::> storage disk removeowner -disk 1.1.23 -data2 true
```

storage disk replace

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 Data 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 Data 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, Data 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
```

storage disk set-foreign-lun

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
```

storage disk set-led

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 Data 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 Data 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 Cluster1* to turn on for 5 minutes:

```
Cluster1::> storage disk set-led -disk Cluster1* -action on -duration 5
```

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 causes the LEDs on all disks whose names match the pattern Cluster1* to turn on for 2 minutes:

```
Cluster1::> storage disk set-led -disk Cluster1* -action on -duration 2
```

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
```

storage disk show

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

| **[-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)

| **[-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

| **[-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 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, Data 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 Data 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.
- Data 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 Data 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".
- Error (if any) while retrieving virtual disk details.

| [-vmdisk-backing-info]

Displays information about the backing disks on certain Data ONTAP-v models:

- Disk name
- Backing disk vendor
- Backing disk model
- Backing disk serial number
- Backing disk device id

| [-foreign] (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

| [-physical-location] (privilege: advanced)

Displays the following information about disks:

- Disk name

- Container type
- Primary path
- Location
- Home node name
- Physical size in human readable units

| [-primary-paths] (privilege: advanced)

Displays the following information about disks:

- Disk Name
- Shelf
- Bay
- Container Type
- Primary Path

| [-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 Data 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 Data 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 (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: `20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000`

[-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.
- array = Logical storage devices backed by storage arrays and used by Data 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|targetasymmap|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 Data ONTAP.
- toobig = The array LUN exceeds the maximum array LUN size that Data ONTAP supports.
- toosmall = The array LUN is less than the minimum array LUN size that Data 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 Data ONTAP.
- netapp = A SAN front-end LUN from one Data ONTAP system that is presented as external storage to another Data 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.

[-hw-minimum-os <text>] - Hardware Minimum Supported Data ONTAP Version

Selects information about disks that have the specified hardware minimum supported Data ONTAP version.

[-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 Data 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.

[-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 (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}] - 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.

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:3ABD5075:00000000:00000000:00000000:00000000:00
000000:00000000
                   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

- storage aggregate show-status

storage disk unfail

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 (privilege: advanced)

This parameter specifies the disk to be unfailed.

`[-s, -spare <true>]` - Make the Disk Spare (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 (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 to become a spare disk:

```
cluster1::*> storage disk unfail -disk 1.1.16 -spare
```

Warning: Failed disk "1.1.16" may have aggregate labels and file system data present. In that case use of this command to release this disk as a spare disk, will lead to permanent loss of any file system data that may be on this disk, and may prevent recovery of a failed aggregate. Are you sure you want to continue with disk unfail? {y|n}:

storage disk updatefirmware

(DEPRECATED) - Update disk firmware

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 Data ONTAP. Use the "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 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

- storage disk firmware update
- storage firmware download

storage disk zerospares

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. Zeroing a disk writes zeros to the entire disk and must be done before a disk can be reused in another aggregate.

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:

```
cluster1::> storage disk zerospares -owner node4
```

storage disk error show

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: `20000000:875D4C32:00000000:00000000:00000000:00000000:00000000:00000000:00000000`

[-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|targetasymmap|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 Data ONTAP.
- toobig = The array LUN exceeds the maximum array LUN size that Data ONTAP supports.
- toosmall = The array LUN is less than the minimum array LUN size that Data 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 Data ONTAP.
- netapp = A SAN front-end LUN from one Data ONTAP system that is presented as external storage to another Data 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.
```

storage disk firmware revert

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 storage disk show`-fields`*firmware-revision* command. The following example displays partial output from the 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 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 (privilege: advanced)

Specifies the disk or disks whose firmware is to be reverted.

| -node {<nodename>|local} - Node Name (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

- storage disk show
- storage firmware download

storage disk firmware show-update-status

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] }

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

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

Selects the node that matches this parameter value.

[-num-waiting-download <integer>] - The Number of Disks Waiting to Download (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) (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) (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 (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 (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 -
```

storage disk firmware update

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 storage disk option modify`-node` *-bkg-firmware -update `*on* command.

You can download the latest firmware on the node by using the 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 storage disk show-fields`*firmware-revision* command. The following example displays partial output from the 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.

Follow the instructions in "How to Update the Firmware for an HA Pair in a Cluster" to ensure that the updating process is successful. Data 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 Backgroud 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 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 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 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 (privilege: advanced)

Specifies the disk or disks whose firmware is to be updated.

| -node {<nodename>|local} - node (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

- storage disk option modify
- storage firmware download
- storage disk show

storage disk option modify

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.
```

storage disk option show

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.
```

storage encryption commands

storage encryption disk destroy

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 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 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 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:

```
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
      xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status] command.

cluster1::>
```

Related Links

- storage encryption disk revert-to-original-state
- storage encryption disk show-status
- storage disk show

storage encryption disk modify

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 and FIPS-compliance protection

parameters of self-encrypting disks (SEDs). The current data AK and FIPS AK of the SED are required to effect changes to the respective AKs and FIPS compliance, and must also be available from the key servers.

The command releases the cluster shell after launching the operation. Monitor the output of the storage encryption disk show-status command for command completion.

> To properly protect data at rest on a SED and place it into compliance with its FIPS certification requirements, set both the Data and FIPS-compliance AKs to a value other than the default manufacture secure ID (MSID), indicated by a key ID with the special value *0x0* . Verify the key IDs by using the storage encryption disk show and `storage encryption disk show-fips` commands.

Parameters

`-disk <disk path name>` - Disk Name

This parameter specifies the name of the 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.

| `[-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 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 authorites and other components as required (other than data) to their default settings. The MSID is required when reverting to a version of Data ONTAP that does not manipulate the FIPS-compliance device components.

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

- storage encryption disk show-status
- storage encryption disk show
- security key-manager query

storage encryption disk revert-to-original-state

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 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 FTPS AK to the default manufacture secure ID (MSID)
- Unlocks the data band
- 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 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 storage disk unfail command in *`advanced`* privilege mode. To do so, you might also need to reestablish ownership of the SED using the 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

- storage disk show
- storage encryption disk show-status
- storage disk unfail
- storage disk assign

storage encryption disk sanitize

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 manufacture secure ID (MSID)
- 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 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 storage disk unfail

command in *advanced* privilege mode. To do so, you might also need to reestablish ownership of the SED using the 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 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
      xref:{relative_path}storage-encryption-disk-show-status.html[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:

```
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
      xref:{relative_path}storage-encryption-disk-show-status.html[storage
encryption disk show-status] command.

cluster1::>
```

Related Links

- storage encryption disk show-status
- storage disk unfail
- storage disk assign
- storage disk show

storage encryption disk show-status

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-sed-support {true|false}] - Node Supports Self-Encrypting Disks

If you specify this parameter, the command displays disk encryption status for the nodes that match this parameter (*true* means the node supports SEDs).

[-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 Data 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.

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 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
```

storage encryption disk show

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 self-encrypting disks (SEDs). By default, the command displays the following information about all SEDS:

- Disk name
- The protection mode of the SED
- The key ID associated with the data authentication key (data AK)

You can use the following parameters together with the `-disk` parameter to narrow the selection of displayed SEDs 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 a SED. If you specify an aggregate name or other container name, only the SEDs in that container are displayed. See the man page for the storage disk show command for a description of the container name. Use the storage aggregate show-status and storage disk show commands to determine which aggregates the SEDs 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 a SED. If you specify a container type, only the SEDs with that container type are displayed. See the man page for the 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 SED requires for authentication with the data-protection authorities in the SED. The special key ID *0x0* indicates that the current data AK of the SED is the default manufacture secure ID (MSID) that is not secret. To properly protect data at rest on the device, modify the data AK using a key ID that is not the MSID. 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 storage encryption disk modify `-data-key-id` *key-id* to protect the data. Use storage encryption disk modify-fips -key-id` *key-id* to place the SED 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 the SED.

[-is-power-on-lock-enabled {true|false}] - Is Power-On Lock Protection Enabled?

This parameter specifies the state of the SED control that determines whether the SED requires authentication with the data AK after a power-cycle. The system enables this control parameter automatically when you use the storage encryption disk modify `-data-key-id` command to set the data AK to a value other than the MSID. Data is protected only when this parameter is *true* and the data AK is not the MSID. 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 SED is in:

- open - data is unprotected; SED is not in FIPS-compliance mode
- data - data is protected; SED is not in FIPS-complance mode
- part - data is unprotected; SED is in FIPS-compliance mode
- full - data is protected; SED is in FIPS-compliance mode

Examples

The following command displays information about all SEDs:

```
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.

The following command displays information about the protection mode and FIPS key ID for all SEDs:

```
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
0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
[...]
```

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.

The following command displays the individual fields for disk 1.10.1:

```
cluster1::> storage encryption disk show -disk 1.10.1
Disk Name: 1.10.1
  Key ID of the Current Data Authentication Key:
0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
  Key ID of the Current FIPS Authentication Key:
0A9C9CFC0000000001000000000000000345CFD1BAD310CA8EDB377D439FB5C9A
           Is Power-On Lock Protection Enabled?: true
Mode of SED Data and FIPS-Compliance Protection: open
```

Related Links

- storage disk show
- storage aggregate show-status
- storage encryption disk modify

storage errors commands

storage errors show

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|targetasymmap|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.
```

storage failover commands

storage failover giveback

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. Some options are available only at the advanced privilege level and higher. Run the 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 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 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 (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 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

- storage failover show-giveback
- storage aggregate show
- storage failover modify

storage failover modify

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 `false` , except for two-node clusters where the default setting is `true` .

[-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 (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 (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) (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 (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) (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 (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) (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) (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 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.

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 managemant 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.

[-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` .

[-aggregate-migration-timeout <integer>] - Aggregate Migration Timeout (secs) (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

- system node reboot

storage failover show-giveback

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.
```

storage failover show-takeover

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.
```

storage failover show

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.

| [-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

| [-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

| [-advanced] (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] (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] (privilege: advanced)

Displays the following information:

- Node name
- Primary mailbox status
- Backup mailbox status

| [-more-options] (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

| [-progress] (privilege: advanced)

Displays the following information:

- Node name
- Maximum resource-table index number
- Current resource-table index number
- Current resource-table entry

| [-timeout] (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] (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 (privilege: advanced)

Selects the nodes that have the specified kill packets setting.

[-partner-aggregates {true|false}] - Control Partner Aggregates (privilege: advanced)

Selects the nodes that have the specified partner aggregates setting.

[-current-index <integer>] - Current Progress Index (privilege: advanced)

Selects the nodes that have the specified current-progress index setting.

[-current-entry <text>] - Current Progress Entry (privilege: advanced)

Selects the nodes that have the specified current-progress entry setting.

[-maximum-index <integer>] - Maximum Progress Index (privilege: advanced)

Selects the nodes that have the specified maximum-progress index setting.

[-pmbox-status <text>,…] - Primary Mailbox Status (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 (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 (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 (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 (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 (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 (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 (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 (privilege: advanced)

Selects the nodes that have the specified local mailbox channel Kernel Time.

[-mbx-sk-cycles-local <integer>] - Local Mailbox CPU Cycles (privilege: advanced)

Selects the nodes that have the specified local mailbox channel CPU Cycle count.

[-ic-abs-time-local <integer>] - Local IC Absolute Time (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 (privilege: advanced)

Selects the nodes that have the specified local Interconnect channel Kernel Time.

[-ic-sk-cycles-local <integer>] - Local IC CPU Cycles (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 (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 (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 (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel Kernel Time.

[-mbx-sk-cycles-partner <integer>] - Partner Mailbox CPU Cycles (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 (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 (privilege: advanced)

Selects the nodes that have the specified partner mailbox channel minor sequence number.

[-ic-abs-time-partner <integer>] - Partner IC Absolute Time (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 (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel Kernel Time.

[-ic-sk-cycles-partner <integer>] - Partner IC CPU Cycles (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 (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 (privilege: advanced)

Selects the nodes that have the specified partner Interconnect channel minor sequence number.

[-local-takeover-info <text>] - Local Takeover Info (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 (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 (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 (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 (privilege: advanced)

Selects the nodes that have the specified fast-timeout configuration setting.

[-slow-timeout <integer>] - Slow Timeout (privilege: advanced)

Selects the nodes that have the specified slow-timeout setting.

[-mailbox-timeout <integer>] - Mailbox Timeout (privilege: advanced)

Selects the nodes that have the specified mailbox-timeout setting.

[-connect-timeout <integer>] - Connect Timeout (privilege: advanced)

Selects the nodes that have the specified connect-timeout setting.

[-operator-timeout <integer>] - Operator Timeout (privilege: advanced)

Selects the nodes that have the specified operator-timeout setting.

[-firmware-timeout <integer>] - Firmware Timeout (privilege: advanced)

Selects the nodes that have the specified firmware-timeout setting.

[-dumpcore-timeout <integer>] - Dumpcore Timeout (privilege: advanced)

Selects the nodes that have the specified dump-core timeout setting.

[-booting-timeout <integer>] - Booting Timeout (privilege: advanced)

Selects the nodes that have the specified booting-timeout setting.

[-transit-timer {true|false}] - Transit Timer Enabled (privilege: advanced)

Selects the nodes that have the specified transit-timer setting.

[-transit-timeout <integer>] - Transit Timeout (privilege: advanced)

Selects the nodes that have the specified transit timeout.

[-firmware-received <integer>] - Firmware Received (privilege: advanced)

Selects the nodes that have the specified firmware-reception time.

[-firmware-received-cycles <integer>] - Firmware Received in CPU Cycles (privilege: advanced)

Selects the nodes that have the specified firmware-reception time in CPU Cycles.

[-booting-received <integer>] - Booting Received (privilege: advanced)

Selects the nodes that have the specified booting-reception time.

[-transit-time <integer>] - Transit Event Time (privilege: advanced)

Selects the nodes whose last failover event occurred at the specified time.

[-pnormal <integer>] - Primary Normal IO Time (privilege: advanced)

Selects the nodes that have the specified normal primary-mailbox I/O time.

[-ptransition <integer>] - Primary Transition IO Time (privilege: advanced)

Selects the nodes that have the specified transitional primary-mailbox I/O time.

[-bnormal <integer>] - Backup Normal IO Time (privilege: advanced)

Selects the nodes that have the specified normal backup-mailbox I/O time.

[-btransition <integer>] - Backup Transition IO Time (privilege: advanced)

Selects the nodes that have the specified transitional backup-mailbox I/O time.

[-logs-unsynced <integer>] - Logs Unsynced Count (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 (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 (privilege: advanced)

Selects the storage-failover configurations that match this parameter value.

[-short-uptime <integer>] - Short Uptime (secs) (privilege: advanced)

Selects the nodes that have the specified short-uptime value.

[-attempts <integer>] - Number of Giveback Attempts (privilege: advanced)

Selects the nodes that have the specified number of giveback attempts.

[-attempts-time <integer>] - Giveback Attempts Period (minutes) (privilege: advanced)

Selects the nodes that have the specified time setting for giveback attempts.

[-propagate {true|false}] - Propagate Status via Mailbox (privilege: advanced)

Selects the nodes that have the specified propagate-status-via-mailbox setting.

[-read-interval <integer>] - Node Status Read Interval (secs) (privilege: advanced)

Selects the nodes that have the specified read interval.

[-write-interval <integer>] - Node Status Write Interval (secs) (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 (privilege: advanced)

Selects the nodes that have the specified firmware state on the local node.

[-local-firmware-progress <integer>] - Local Firmware Progress Counter (privilege: advanced)

Selects the nodes that have the specified firmware progress counter for the local node.

[-partner-firmware-state <text>] - Partner Firmware State (privilege: advanced)

Selects the nodes that have the specified firmware state of the partner node.

[-partner-firmware-progress <integer>] - Partner Firmware Progress Counter (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 (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 (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 (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 (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 (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 (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) (privilege: advanced)

Selects the nodes that have the specified aggregate migration timeout.

[-is-mirror-enabled {true|false}] - Is NVRAM Mirroring Enabled (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 (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 (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 (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 (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 (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 (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 (privilege: advanced)

Selects the nodes that have the specified total system uptime, in milliseconds.

[-current-time <integer>] - Current System Time (privilege: advanced)

Selects the nodes that have the specified current time on the node.

[-fm-takeover-state <FM Takeover/Giveback Transition>] - FM Takeover State (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 (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 (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 occured during the boot process
- TAKEOVER_OPERATOR_EXP - Takeover occured after the operator timeout expired
- TAKEOVER_POST_FAILED - Takeover occured 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.
```

storage failover takeover

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.

  If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass-optimization` option is set to false.

  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.

  If this option is specified, negotiated takeover optimization is bypassed even if the `-bypass-optimization` option is set to false.

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

If this is an operator-initiated planned takeover, this parameter specifies whether the takeover optimization is bypassed. This parameter defaults to false.

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.

If this parameter is specified, negotiated takeover optimization is bypassed even if the `-bypass-optimization` parameter is set to false.

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.

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
```

storage failover hwassist show

Display hwassist status

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

Description

The `storage failover hwassist show` command displays information about hardware assisted takeover configurations. By default, the command displays the following information:

- Node name.
- Partner node name.
- Whether hardware assisted takeover is enabled.
- IP address on which the local node receives hardware assist alerts.
- Port on which local node receives hardware assist alerts.
- Hardware assist monitor status.
- If the monitor is inactive, the reason it is inactive.
- If the monitor is inactive, the corrective action to make it active.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 configurations that match this parameter value.

[-partner-name {<nodename>|local}] - Name of the Partner Node

Selects the hwassist configurations that match this parameter value.

[-enabled {true|false}] - Local Hardware Assist Enabled

Selects the hwassist configurations that match this parameter value.

[-local-status <text>] - Local Node's Hwassist Status

Selects the hwassist configurations that match this parameter value (*`active`* or *`inactive`*).

[-local-ip <text>] - IP Address on Which Local Node is Listening

Selects the hwassist configurations that match this parameter value.

[-local-port <integer>] - Port on Which Local Node is Listening

Selects the hwassist configurations that match this parameter value.

[-local-inactive <text>] - Local Node's Hwassist Inactive Status Reason

Selects the hwassist configurations that match this parameter value.

[-local-action <text>] - Corrective Action on Local Node

Selects the hwassist configurations that match this parameter value.

Examples

The following example displays the hardware assist information for the local node and its partner:

```
cluster1::> storage failover hwassist show
    Node
    -----------------
    ha1
                                Partner : ha2
                      Hwassist Enabled : true
                           Hwassist IP : 10.225.248.19
                         Hwassist Port : 4444
                        Monitor Status : active
                       Inactive Reason : -
                     Corrective Action : -
    ha2
                                Partner : ha1
                  Hwassist Enabled : true
                       Hwassist IP : 10.225.248.21
                     Hwassist Port : 4444
                    Monitor Status : active
                   Inactive Reason : -
                 Corrective Action : -
```

storage failover hwassist test

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.
```

storage failover hwassist stats clear

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
```

storage failover hwassist stats show

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.

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  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
    ------------ -------------------- ------ ---------
--------------------
    -
```

storage failover internal-options show

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] (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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects configuration information for the specified node.

[-auto-giveback {true|false}] - Auto Giveback Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified automatic giveback setting.

[-check-partner {true|false}] - Check Partner Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified partner-checking setting.

[-detection-time <integer>] - Takeover Detection Time (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified takeover detection time setting.

[-onfailure {true|false}] - Takeover on Failure Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-failure setting.

[-onpanic {true|false}] - Takeover on Panic Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-panic setting.

[-onshort-uptime {true|false}] - Takeover on Short Uptime Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-short-uptime setting.

[-short-uptime <integer>] - Short Uptime (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-short-uptime time setting.

[-attempts <integer>] - Number of Giveback Attempts (privilege: advanced)

Selects configuration information for nodes that have the specified number of giveback attempts setting.

[-attempts-time <integer>] - Giveback Attempts Minutes (privilege: advanced)

Selects configuration information for nodes that have the specified giveback attempt time setting.

[-propagate {true|false}] - Propagate Status via Mailbox (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) (privilege: advanced)

Selects configuration information for nodes that have the specified status read interval setting.

[-write-interval <integer>] - Node Status Write Interval (secs) (privilege: advanced)

Selects configuration information for nodes that have the specified status write interval setting.

[-onreboot {true|false}] - Takeover on Reboot Enabled (privilege: advanced)

Selects configuration information for nodes that have the specified takeover-on-reboot setting.

[-delay-seconds <integer>] - Delay Before Auto Giveback (secs) (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 (privilege: advanced)

Selects configuration information for nodes that have the specified hardware-assisted takeover setting.

[-hwassist-partner-ip <text>] - Partner's Hwassist IP (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 (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) (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 (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 (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 (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
```

storage failover mailbox-disk show

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] }

If this parameter is used, the command displays detailed information about all entries.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the mailbox disks that are associated with the specified node.

[-location {local|partner|tertiary}] - Mailbox Owner (privilege: advanced)

Selects the mailbox disks that have the specified relationship to the node.

[-diskindex <integer>] - Mailbox Disk Index (privilege: advanced)

Selects the mailbox disk that has the specified index number.

[-diskname <text>] - Mailbox Disk Name (privilege: advanced)

Selects the mailbox disks that match the specified disk name.

[-diskuuid <text>] - Mailbox Disk UUID (privilege: advanced)

Selects the mailbox disks that match the specified UUID.

[-physical-location {local|partner|mediator}] - Mailbox Disk Physical Location (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 (privilege: advanced)

Selects the mailbox disks that match the specified location-id.

[-location-name <text>] - Mailbox Disk Location (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:[...]
```

storage failover progress-table show

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] }

If this parameter is used, the command displays detailed information about all entries.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects the status information for the specified node.

[-index <integer>] - Resource Table Index (privilege: advanced)

Selects the status information for the specified index number.

[-entryname <text>] - Resource Table Entry Name (privilege: advanced)

Selects the status information for the specified entry name.

[-state <text>] - Resource Table Entry State (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 (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 (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
```

storage firmware commands

storage firmware download

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 storage disk firmware update command to install downloaded disk firmware.

Use the storage shelf firmware update command to install downloaded shelf module firmware.

Use the system node run` local storage download acp` command to install downloaded ACP processor firmware.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This specifies the node to which the firmware is to be downloaded.

`-package-url <text>` - Package URL (privilege: advanced)

This specifies the path to the firmware package.

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

- storage disk firmware update
- storage shelf firmware update
- system node run

storage firmware acp delete

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 (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
```

storage firmware acp rename

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 (privilege: advanced)

This parameter specifies the firmware file to rename.

`-newname <text>` - New Filename (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
```

storage firmware acp show

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.
```

storage firmware disk delete

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 (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
```

storage firmware disk rename

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 (privilege: advanced)

This parameter specifies the firmware file to rename.

`-newname <text>` - New Filename (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
```

storage firmware disk show

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.
```

storage firmware shelf delete

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 (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
```

storage firmware shelf rename

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 (privilege: advanced)

This parameter specifies the firmware file to rename.

`-newname <text>` - New Filename (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
```

storage firmware shelf show

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.
```

storage iscsi-initiator commands

storage iscsi-initiator add-target

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 (privilege: advanced)

Specifies the name of the Data ONTAP node to which the iSCSI target will be added.

`-label <text>` - User Defined Identifier (privilege: advanced)

Specifies a label for the target to be added.

`-target-type {external|mailbox|partner|dr_auxiliary|dr_partner}` - Target Type (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 (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 (privilege: advanced)

Specifies the iSCSI target name such as an IQN (iSCSI qualified name).

`[-status-admin {down|up}]` - Administrative Status (default: up) (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
```

storage iscsi-initiator connect

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 (privilege: advanced)

Specifies the name of the Data ONTAP node to which the iSCSI target will be connected.

`[-target-type {external|mailbox|partner|dr_auxiliary|dr_partner}]` - Target Type (privilege: advanced)

Selects targets with the specified target type.

`-label <text>` - User Defined Identifier (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
```

storage iscsi-initiator disconnect

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 (privilege: advanced)

Specifies the name of the Data ONTAP node from which the iSCSI target will be disconnected.

`[-target-type {external|mailbox|partner|dr_auxiliary|dr_partner}]` - Target Type (privilege: advanced)

Selects targets with the specified target type.

`-label <text>` - User Defined Identifier (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
```

storage iscsi-initiator remove-target

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 (privilege: advanced)

Specifies the name of the Data ONTAP node from which the iSCSI target will be removed.

`[-target-type {external|mailbox|partner|dr_auxiliary|dr_partner}]` - Target Type (privilege: advanced)

Selects targets with the specified target type.

`-label <text>` - User Defined Identifier (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
```

storage iscsi-initiator show

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 Data 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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Represents the name of the Data 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|dr_auxiliary|dr_partner}] - Target Type (privilege: advanced)

Selects targets with the specified target type.

[-label <text>] - User Defined Identifier (privilege: advanced)

Selects targets with the specified label.

[-target-portal <text>] - Target Portal (privilege: advanced)

Selects targets with the specified portal.

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

Selects targets with the specified target name.

[-status-admin {down|up}] - Administrative Status (privilege: advanced)

Selects targets with the specified administrative status.

[-status-oper {down|up}] - Operational Status (privilege: advanced)

Selects targets with the specified operational status.

[-failure-reason <text>] - Failure Reason (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.
```

storage load commands

storage load balance

Balance storage I/O across controller's initiator ports

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

Description

This command is obsolete. I/O load is balanced automatically every five minutes.

Parameters

`-node {<nodename>|local}` - Node to balance on

The name of the clustered node for which information is being displayed.

Examples

This command has no effect.

storage load show

Display I/O statistics to array LUNs, grouped by initiator port.

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

Description

This command is obsolete. The `storage load show` command displays the load distribution of I/O on 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.

| [-switch]

The switch parameter adds switch information to the 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-port <text>] - Initiator Port

The initiator port of the array LUN for which I/O stats are being displayed.

[-wwpn <text>] - Target Port WWPN

The World Wide Port Name of the array LUN for which I/O stats are being displayed.

[-serialnumber <text>] - Serial Number

The serial number of the array LUN for which I/O stats are being displayed.

[-lun <integer>] - LUN

The array LUN for which I/O stats are being displayed.

[-pct-io <text>] - %I/O

Percent of I/O bandwidth consumed by this array LUN.

[-io-blocks <integer>] - I/O (blocks)

Number of I/O blocks transferred.

[-switch-port <text>] - Switch Port

The initiator side switch port for this array LUN.

[-target-side-switch-port <text>] - Target Side Switch Port

The target side switch port for this array LUN.

Examples

```
cluster1::> storage load show -switch
Initiator port: 0a connected to vnbr3850s4:7.
                                                                    Target Side
LUN Serial #                              Target Port               Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  1 D600020C00D3                      50060e80004291c0        vnbr3850s5:12
0            0
  2 D600020C00D4                      50060e80004291c0        vnbr3850s5:12
100           21
  5 D600020C00EF                      50060e80004291c0        vnbr3850s5:12
0            0
Initiator port: 0c connected to vnci9124s54:1-11.
                                                                    Target Side
LUN Serial #                              Target Port               Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  3 D600020C00D9                      50060e80004291c2     vnci9124s54:1-22
42            8
  4 D600020C00DA                      50060e80004291c2     vnci9124s54:1-22
36            7
  6 D600020C00F0                      50060e80004291c2     vnci9124s54:1-22
15            3
Initiator port: 0a connected to vnbr3850s4:8.
                                                                    Target Side
LUN Serial #                              Target Port               Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  2 D600020C00D4                      50060e80004291c0        vnbr3850s5:12
0            0
Initiator port: 0a connected to vnbr3850s4:8.
                                                                    Target Side
LUN Serial #                              Target Port               Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  5 D600020C00EF                      50060e80004291c0        vnbr3850s5:12
0            0
  6 D600020C00F0                      50060e80004291c0        vnbr3850s5:12
100           31
Initiator port: 0c connected to vnci9124s54:1-12.
```

```
                                                                  Target Side
LUN Serial #                              Target Port             Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  1 D600020C00D3                     50060e80004291c2     vnci9124s54:1-22
0            0
Initiator port: 0c connected to vnci9124s54:1-12.
                                                                  Target Side
LUN Serial #                              Target Port             Switch Port
%I/O I/O (blocks)
--- ------------------------------ ---------------- --------------------
------ ------------
  3 D600020C00D9                     50060e80004291c2     vnci9124s54:1-22
42           3
  4 D600020C00DA                     50060e80004291c2     vnci9124s54:1-22
42           3
12 entries were displayed.
```

storage path commands

storage path quiesce

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 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, Data 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

- storage path resume

storage path resume

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 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, Data 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

- storage path quiesce

storage path show-by-initiator

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 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

- storage path show

storage path show

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.
```

storage pool commands

storage pool add

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 or usable capacity of that aggregate is increased depending on whether it is a Flash Pool or an All-Flash aggregate .

If capacity from a storage pool is already provisioned into a Flash Pool aggregate, the same storage pool cannot be used to provision an All-Flash aggregate and vice-versa.

For provisioning storage pool capacity into All-Flash aggregates, the Vserver option *`raid.storagepool.data.enable`* must be set to `true` . The storage pool data enabled mode of operation is not currently supported by OnCommand management software.

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 or usable capacity by 186.2 GB. If a different allocation is desired, create a new storage pool using the storage pool create command.

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 (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

- storage pool create

storage pool create

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, Data 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 storage pool reassign command.

After the storage pool is created, its allocation units can be provisioned into Flash Pool or All-Flash aggregates using the storage aggregate add-disks command and the `-storage-pool` parameter.

If capacity from a storage pool is already provisioned into a Flash Pool aggregate, the same storage pool cannot be used to provision an All-Flash aggregate and vice versa.

For provisioning storage pool capacity into All-Flash aggregates, the vserver option *`raid.storagepool.data.enable`* must be set to `true` . The storage pool data enabled mode of operation is not currently supported by OnCommand management software.

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

- storage pool reassign
- storage aggregate add-disks

storage pool delete

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
```

storage pool reassign

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.
```

storage pool show-aggregate

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 or All-Flash aggregates respectively .

[-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
```

storage pool show-available-capacity

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 or All-Flash aggregates. Allocation units might be provisioned into aggregates using the storage aggregate add-disks command and the `-storage-pool` parameter.

> 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, Data 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 Data 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

- storage aggregate add-disks

storage pool show-disks

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}] - Disk Type

Selects the storage pools with the disks that match this parameter value. Only the *SSD* type is supported for this version of Data 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
```

storage pool show

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

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 Data 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: -
```

storage port commands

storage port disable

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 (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
```

storage port enable

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
```

storage port rescan

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.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the node on which the port resides.

`-port <text>` - Port (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
```

storage port reset-device

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.
```

Parameters

-node {<nodename>|local} - Node (privilege: advanced)
Use this parameter to specify the node on which the port resides.

-port <text> - Port (privilege: advanced)
Use this parameter to specify the port used to reset the device.

{ -shelf-name <text> - Shelf Name (privilege: advanced)
Use this parameter to specify the shelf where the device resides.

-bay-id <integer> - Bay ID (privilege: advanced)
Use this parameter to specify the bay where the device resides.

| -loop-id <integer> - Loop ID (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
```

storage port reset

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.

Parameters

-node {<nodename>|local} - Node (privilege: advanced)
Use this parameter to specify the node on which the port resides.

-port <text> - Port (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
```

storage port show

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}] - Port Type

Selects the ports of the specified type.

[-port-speed {0|1|1.5|2|3|4|6|8|10|12|16|32}] - Port Speed

Selects the ports with the specified speed.

[-state {enabled|disabled}] - Port State

Selects the ports with the specified state.

[-status {unknown|online|online-degraded|offline|link-down}] - 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.

[-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.

[-part-number <text>] - Part Number

Selects the ports with the specified part number.

[-date-code <text>] - Date Code

Selects the ports with the specified date code.

[-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-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|3|4|6|8|10|12|16|32}] - Phy Speed

Selects the ports that have phys with the specified speed.

[-error-type {unknown|online|online-degraded|offline|link-down}] - 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.

[-connector-type {QSFP|QSFP+|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.

Examples

The following example displays information about all storage ports in the cluster:

```
cluster1::> storage port show
Speed
Node               Port Type (Gb/s) State    Status
------------------ ---- ---- ------ -------- ---------------
node1
                   0a   SAS       0 disabled offline
                   0b   SAS       6 enabled  online
                   1a   FC        0 disabled offline
                   1b   FC        2 enabled  online
node2
                   0a   SAS       6 enabled  online
                   0b   SAS       6 enabled  online-degraded
                   1a   FC        2 enabled  online
                   1b   FC        2 enabled  online
```

The following example displays detailed information about port 0a on node node1:

```
cluster1::> storage port show -node node1 -port 0b
Node: node1
                   Port: 0b
            Description: SAS Host Adapter 0b (PMC-Sierra PM8001 rev. C)
      Firmware Revision: 01.12.06.00
               Base WWN: 50:0a:09:80:00:82:47:b4
         Connector Type: qsfp
       Connector Vendor: Molex Inc.
  Connector Part Number: 112-00177+A0
   Connector Technology: passive_copper
Connector Serial Number: 017920547
           Cable Length: 2m
   Cable End Identifier: end_1
       Cable Identifier: 500a09800082847f-500a0980008247b4
             Port Speed: 6 Gb/s
             Port State: enabled
            Port Status: online
              Phy State: [0]  enabled,  online,  6 Gb/s
                         [1]  enabled,  online,  6 Gb/s
                         [2]  enabled,  online,  6 Gb/s
                         [3]  enabled,  online,  6 Gb/s
```

storage raid-options commands

storage raid-options modify

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 storage raid-options show command.

`[-value <text>]` - Option Value

This parameter specifies the value of the selected RAID option.

Related Links

- storage raid-options show

storage raid-options show

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.

[-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*
Node     Option                                Value        Constraint
-------- ------------------------------------ ------------ -----------
node1    raid.scrub.perf_impact                low          only_one
node1    raid.scrub.schedule                                none
3 entries were displayed.
```

storage shelf commands

storage shelf show

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 which 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 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.

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

Displays information only about the storage shelves that are attached to the node you specify.

[-shelf <text>] - Shelf Name

Displays information only about the storage shelves that match the names 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|iom12b}] - Shelf Module Type

Displays information only about the storage shelves that match the module-type you specify.

[-connection-type {unknown|fc|sas}] - 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 a 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 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}] - 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_to_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_drv_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_drive|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_wrd_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_dis_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_burst|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_dis_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown}] - 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_to_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_drv_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_drive|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_wrd_rate|term|phy_dis_clk_fault|phy_dis_crc_err|phy_dis_crc_err_burst|phy_dis_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_burst|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_dis_phy_unused|phy_ena|phy_ena_not_attach|phy_ena_unknown|phy_unknown}] - 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.

[-error-type {unknown|acpp|bay|configuration|current|disk|internal|fan|module|port|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).
```

storage shelf acp configure

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.

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
```

storage shelf acp show

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 specifed 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:

```
cluster1::> storage shelf acp show
     Node                Channel                Connectivity
        ----------------  --------------------   ----------------------
     stor-8020-1         in-band                active
     stor-8020-2         in-band                active
     stor-8060-1         out-of-band            full-connectivity
     stor-8060-2         out-of-band            full-connectivity
     4 entries were displayed.
```

The following example displays the -instance output of the storage acp show. More details on the connectivity and configuration can be seen here.

```
cluster1::> storage shelf acp show -instance
Node: stor-8020-1
           Channel: out-of-band
        Enable Status: enabled
                  Port: e0P
           IP Address: 192.168.1.74
               Subnet: 192.168.0.1
              Netmask: 255.255.252.0
    Connection Status: full-connectivity
Node: stor-8020-1
           Channel: in-band
        Enable Status: enabled
    Connection Status: active
2 entries were displayed.
```

storage shelf acp module show

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}] - 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.

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
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
```

```
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
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
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
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
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
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

8 entries were displayed.
```

storage shelf drawer show-phy

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::>
```

storage shelf drawer show-slot

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::>
```

storage shelf drawer show

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::>
```

storage shelf firmware show-update-status

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] }

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

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

Selects the node that matches this parameter value.

[-update-status {running|idle}] - Disk Shelf Firmware Update Status (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 (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
```

storage shelf firmware update

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 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 (privilege: advanced)

This specifies the name of the shelf whose firmware is to be updated.

| [-module-type {atfcx|esh4|iom3|iom6|iom6e|iom12|iom12e|iom12f|iom12b}] - Shelf Module Type (privilege: advanced) }

Update the firmware on the shelves that match the module-type you specify.

[-refresh <true>] - Refresh (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

- storage firmware download

storage shelf location-led modify

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 -node node1 -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 -node node1 -shelf-name 1.0
-led-status off

Info: Shelf locate request successful for shelf "1.0".
```

storage shelf location-led show

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.
```

storage switch commands

storage switch add

Add a back-end switch for monitoring

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

Description

The `storage switch add` command enables you to add 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 switch that is added for monitoring.

[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}] - Supported SNMP Version

This parameter specifies the SNMP version that Data ONTAP uses to communicate with the back-end switch that is added for monitoring. The default SNMP version is SNMPv2c.

{ [-snmp-community <text>] - (DEPRECATED)-SNMPv2c Community or SNMPv3 Username

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-snmp-community-or-username` instead.

This parameter specifies the SNMPv2c community set or SNMPv3 username on the switch that is added for monitoring.

| [-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? (privilege: advanced)

If specified, the `storage switch 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::> storage switch add -address 10.226.197.34 -snmp-community-or
-username public
cluster1::> storage switch 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. Data ONTAP uses SNMPv3 and 'snmpuser1' username to communicate with this switch.

```
cluster1::> storage switch add -address 10.228.56.208 -snmp-version SNMPv3
-snmp-community-or-username snmpuser1 -blades 3,4
```

storage switch modify

Modify information about a back-end switch's configuration

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

Description

The `storage switch modify` enables you to modify certain parameters for identifying and accessing the FC back-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 Data 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 <text>] - (DEPRECATED)-SNMPv2c Community or SNMPv3 Username

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-snmp-community-or-username` instead.

This parameter specifies the SNMPv2c community set or SNMPv3 username on 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::> storage switch 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::> storage switch 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::> storage switch modify -switch-name Brocade6505 -switch
-ipaddress 10.226.197.34 -snmp-version SNMPv3 -snmp-community-or-username
snmpuser1
    cluster1::>
```

storage switch refresh

Refresh storage switch info

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

Description

The `storage switch refresh` command triggers a refresh of the SNMP data for the MetroCluster 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 storage switch add and storage bridge add commands respectectively.

Examples

The following command triggers a refresh for the SNMP data:

```
cluster1::*> storage switch refresh
cluster1::*>
```

Related Links

- storage switch add
- storage bridge add

storage switch remove

Remove a back-end switch from monitoring

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

Description

The `storage switch remove` enables you to remove FC back-end switches that were previously added for SNMP monitoring.

Parameters

`-switch-name <text>` - FC Switch Name

This parameter specifies the name of the back-end switch added for monitoring.

Examples

The following command removes 'Cisco_10.226.197.34' switch from monitoring:

```
cluster1::> storage switch 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::> storage switch remove -switch-name Cisco_10.226.197.34
cluster1::> storage switch 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::>
```

storage switch show

Display switch information

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

Description

The `storage switch show` command displays information about all the storage switches in the MetroCluster configuration. The back-end switches must have been previously added for monitoring using the storage switch add command. If no parameters are specified, the default command displays the following information about the storage switches:

- Switch
- Symbolic Name
- Vendor
- Model
- Switch WWN
- Is Monitored
- Monitor Status

To display detailed profile information about a single storage switch, use the `-switch-name` parameter.

Parameters

{ [-fields <fieldname>,…]

Displays the specified fields for all the storage switches, in column style output.

| [-connectivity]

Displays the following details about the connectivity from the storage 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 storage 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 storage switch:

- Fan name
- Fan speed in rotations per minute (RPM)
- Fan operational status
- Temperature sensor name
- Temperature sensor reading in Celsius ©
- Temperature sensor status

| [-error]

Displays the errors related to the storage switch.

| [-port]

Displays the following details about the storage 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 storage switch power supplies:

- Power supply name
- Power supply serial number
- Power supply operational status

| [-san-config]

Displays the following details about the Virutal Storage Area Networks (VSAN) and Zones of the storage 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 storage switch ports Small Formfactor 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 storage 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 storage switches in the system. If a storage switch is specified, then this parameter displays the same detailed information for the storage switch you specify as does the -switch-name parameter.

[-switch-name <text>] - FC Switch Name

Displays information only about the storage switches that match the name you specify.

[-switch-wwn <text>] - Switch World Wide Name

Displays information only about the storage switches that match the switch wwn you specify.

[-switch-symbolic-name <text>] - Switch Symbolic Name

Displays information only about the storage switches that match the switch symbolic name you specify.

[-switch-fabric-name <text>] - Fabric Name

Displays information only about the storage switches that match the switch fabric you specify.

[-domain-id <integer>] - Switch Domain ID

Displays information only about the storage switches that match the switch domain id you specify.

[-switch-role {unknown|primary|subordinate}] - Switch Role in Fabric

Displays information only about the storage switches that match the switch role you specify.

[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}] - SNMP Version

Displays information only about the storage switches that match the switch SNMP version you specify.

[-switch-model <text>] - Switch Model

Displays information only about the storage switches that match the switch model you specify.

[-switch-vendor {unknown|Brocade|Cisco}] - Switch Vendor

Displays information only about the storage switches that match the switch vendor you specify.

[-fw-version <text>] - Switch Firmware Version

Displays information only about the storage switches that match the switch firmware version you specify.

[-serial-number <text>] - Switch Serial Number

Displays information only about the storage switches that match the switch serial number you specify.

[-switch-ipaddress <IP Address>] - Switch IP Address

Displays information only about the storage switches that match the switch IP address you specify.

[-switch-status {unknown|ok|error}] - Switch Status

Displays information only about the storage switches that match the switch status you specify.

[-snmp-community <text>] - (DEPRECATED)-SNMPv2c Community or SNMPv3 Username

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-snmp-community-or-username` instead.

Displays information only about the storage switches that match the switch SNMPv2c community or SNMPv3 username you specify.

[-snmp-community-or-username <text>] - SNMPv2c Community or SNMPv3 Username

Displays information only about the storage 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 storage 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 storage switches that match the switch monitoring value you specify.

[-blades <integer>,…] - Director-Class Switch Blades to Monitor

Displays information only about the storage switches that match the blade value you specify.

[-engine-id <Hex String>] - Engine ID of SNMPv3 Capable Switch

Displays information only about the storage switches that match the SNMPv3 engine-id you specify.

[-psu-name-list <text>,…] - Switch Power Supply Name List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage switches that have the temperature sensors with the names you specify.

[-temp-sensor-reading-list <integer>,…] - Switch Temperature Sensor Reading © List

Displays information only about the storage 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 storage 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 storage 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 storage switches that match the fans with the names you specify.

[-fan-rpm-list <integer>,…] - Switch Fan Speed (RPM) List

Displays information only about the storage 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 storage 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 storage 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 storage switches that have the VSANs with the indexes you specify.

[-vsan-name-list <text>,…] - Switch VSAN Name List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage switches that have the VSANs with the memeber switch port names you specify.

[-vsan-id-list <integer>,…] - Zone VSAN ID List

Displays information only about the storage switches that have the VSANs with the IDs you specify.

[-zone-name-list <text>,…] - Switch Zone Name List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage switches that have the ports with the WWNs you specify.

[-port-name-list <text>,…] - Switch Port Name List

Displays information only about the storage 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 storage 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 storage 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}] - Switch Port Mode List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage switches that have the errors you specify.

[-conn-switch-port-name-list <text>,…] - Switch Port Name List

Displays information only about the storage 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}] - Switch Port Operating Mode List

Displays information only about the storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage 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 storage switches that have the ports with the speeds you specify.

[-node-name-list <nodename>,…] - Node Name List

Displays information only about the storage switches that are connected to the nodes you specify.

[-adapter-name-list <text>,…] - Node Adapter Name List

Displays information only about the storage switches that are connected to the adapters you specify.

[-adapter-port-name-list <text>,…] - Node Adapter Port Name List

Displays information only about the storage 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 storage 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 storage 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 storage switches that match the names you specify.

[-domain-id-list <integer>,…] - Switch Domain ID List

Displays information only about the storage switches that match the domain ids you specify.

[-wwn-list <text>,…] - Switch WWN List

Displays information only about the storage switches that match the switch WWNs you specify.

[-role-list {unknown|primary|subordinate}] - Switch Role in Fabric List

Displays information only about the storage switches that match the switch roles you specify.

[-address-list <IP Address>,…] - Switch IP Address List

Displays information only about the storage switches that match the switch IP addresses you specify.

Examples

The following example displays information about all storage switches:

```
cluster::>storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::>storage switch 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 (C) 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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 storage switches:

```
cluster::> storage switch 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

- storage switch add

storage tape commands

storage tape offline

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
```

storage tape online

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
```

storage tape position

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
```

storage tape reset

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 (privilege: advanced)

Use this parameter to specify the node to which the tape drive is attached.

-device-id <text> - Tape Drive Device ID (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
```

storage tape show-errors

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: -
```

storage tape show-media-changer

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
```

storage tape show-supported-status

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 Data 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 Data 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
```

storage tape show-tape-drive

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
```

storage tape show

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.

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
```

storage tape trace

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 disbled.

`[-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
```

storage tape alias clear

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
```

storage tape alias set

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 90 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].
```

storage tape alias show

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.

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
```

storage tape config-file delete

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 (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
```

storage tape config-file get

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
```

storage tape config-file show

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
```

storage tape library config show

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::>
```

storage tape library path show-by-initiator

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.
```

storage tape library path show

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.
```

storage tape load-balance modify

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
```

storage tape load-balance show

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.
```

system commands

system chassis commands

system chassis show

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
```

system chassis fru show

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}] - 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
```

system cluster-switch commands

system cluster-switch create

Add information about a cluster switch or management switch

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

Description

The `system cluster-switch create` command adds information about a cluster switch or management switch. The cluster switch health monitor uses this information to monitor the health of the switch.

Use this command if ONTAP cannot automatically discover a cluster or management switch. ONTAP relies on the Cisco Discovery Protocol (CDP) to discover the switches. CDP is always enabled on all cluster ports of a node by default, disabled on all non-cluster ports of a node. If the CDP is also enabled on your cluster switches, they will be automatically discovered.

If you want ONTAP to discover and monitor management switches, the CDP must be enabled on non-cluster ports. To verify whether the CDP is enabled or disabled, use the command system node run`-node` `<node_name>`-command`options cdpd.enable`.

Use the system cluster-switch show command to identify switches that the cluster switch health monitor is monitoring.

Parameters

`-device <text>` - Device Name

Specifies the device name of the switch that you want to monitor. Data ONTAP uses the device name of the switch to identify the SNMP agent with which it wants to communicate.

`-address <IP Address>` - IP Address

Specifies the IP address of switch's management interface.

`-snmp-version {SNMPv1|SNMPv2c|SNMPv3}` - SNMP Version

Specifies the SNMP version that Data ONTAP uses to communicate with the switch. The default is SNMPv2c.

{ `-community <text>` - DEPRECATED-Community String or SNMPv3 Username

> This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-community-or-username` instead.

Specifies the community string for SNMPv2 authentication or SNMPv3 user name for SNMPv3 security. The default community string for SNMPv2 authentication is cshm1!.

| `-community-or-username <text>` - Community String or SNMPv3 Username }

Specifies the community string for SNMPv2 authentication or SNMPv3 user name for SNMPv3 security. The default community string for SNMPv2 authentication is cshm1!.

`-model {NX5010|NX5020|CAT2960|OTHER|NX5596|CN1610|CN1601|NX3132|OT5548|NX3132V|OT9332|NX3132XL}` - Model Number

Specifies the model number of the switch. You should not set this parameter to OTHER. Data ONTAP does not monitor switches that match this value. Data ONTAP sets this parameter to OTHER if a switch that it automatically discovers is not supported for health monitoring.

`-type {cluster-network|management-network}` - Switch Network

Specifies the switch type.

[-is-monitoring-enabled-admin {true|false}] - Enable Switch Monitoring

Specifies the switch admin monitoring status.

Examples

```
cluster1::> system cluster-switch create -device SwitchA -address 1.2.3.4
-snmp-version SNMPv2c -community-or-username cshm1! -model NX55596 -type
cluster-network
```

Creates a new switch configuration for a switch named SwitchA.

```
cluster2::> system cluster-switch create -device SwitchB -address 5.6.7.8
-snmp-version SNMPv3 -community-or-username snmpv3u1 -model CN1601 -type
management-network
```

Creates a new switch configuration for a switch named SwitchB.

Related Links

- system node run
- system cluster-switch show

system cluster-switch delete

Delete information about a cluster switch or management switch

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

Description

The `system cluster-switch delete` command disables switch health monitoring for a cluster or management switch.

Parameters

-device <text> - Device Name

Specifies the name of the switch.

[-force <true>] - Force Delete (privilege: advanced)

Specifies if force delete or not.

Examples

```
cluster1::> system cluster-switch delete -device SwitchA
```

Disables monitoring for the switch named SwitchA.

```
cluster1::> system cluster-switch delete -device SwitchA -force
```

Forcefully disables monitoring for the switch named SwitchA. (privilege: advanced)

system cluster-switch modify

Modify information about a switch's configuration

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

Description

The ` system cluster-switch modify ` command modifies information about a cluster switch or management switch. The cluster switch health monitor uses this information to monitor the switch.

Parameters

`-device <text>` - Device Name

Specifies the device name of switch that you want to monitor.

`[-address <IP Address>]` - IP Address

Specifies the IP address of switch's management interface.

`[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}]` - SNMP Version

Specifies the SNMP version that Data ONTAP uses to communicate with the switch. The default is SNMPv2c.

{ `[-community <text>]` - DEPRECATED-Community String or SNMPv3 Username

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-community-or-username` instead.

Specifies the community string for SNMPv2 authentication or SNMPv3 username for SNMPv3 security.

| `[-community-or-username <text>]` - Community String or SNMPv3 Username }

Specifies the community string for SNMPv2 authentication or SNMPv3 username for SNMPv3 security.

`[-type {cluster-network|management-network}]` - Switch Network

Specifies the switch type.

`[-is-monitoring-enabled-admin {true|false}]` - Enable Switch Monitoring

Specifies the switch admin monitoring status.

Examples

```
cluster1::> system cluster-switch modify -device SwitchA -address 2.3.4.5
```

Modifies the IP address for the switch named SwitchA.

```
cluster1::> system cluster-switch modify -device SwitchB -snmp-version
SNMPv3 -community-or-username snmpv3u1
```

Modifies the SNMP parameters for the switch named SwitchB.

system cluster-switch prepare-to-downgrade

Remove unsupported switches in preparation for downgrade

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

Description

The `system cluster-switch prepare-to-downgrade` command changes switch information, so that it is compatible with older versions of ONTAP. When executed, it removes cluster switch entries that are not supported in versions earlier than ONTAP 9.1.

Examples

cluster1::> system cluster-switch prepare-to-downgrade

system cluster-switch show-all

Displays the list of switches that were added and deleted

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

Description

The `system cluster-switch show-all` command displays configuration details for discovered monitored cluster switches and management switches, including switches that are user-deleted. From the list of deleted switches, you can delete a switch permanently from the database to re-enable automatic discovery of that switch.

Parameters

{ [-fields <fieldname>,…]
Selects the fields that have the specified name.

| [-instance] }
Selects detailed information for all the switches.

[-device <text>] - Device Name (privilege: advanced)
Selects the switches that match the specified device name.

[-address <IP Address>] - IP Address (privilege: advanced)
Selects the switches that match the specified IP address.

[-snmp-version {SNMPv1|SNMPv2c|SNMPv3}] - SNMP Version (privilege: advanced)

Selects the switches that match the specified SNMP version.

[-community <text>] - DEPRECATED-Community String or SNMPv3 Username (privilege: advanced)

> This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-community-or-username` instead.

Selects the switches that match the specified community string or SNMPv3 username.

[-community-or-username <text>] - Community String or SNMPv3 Username (privilege: advanced)

Selects the switches that match the specified community string or SNMPv3 username.

[-discovered {true|false}] - Is Discovered (privilege: advanced)

Selects the switches that match the specified discovery setting.

[-model {NX5010|NX5020|CAT2960|OTHER|NX5596|CN1610|CN1601|NX3132|OT5548|NX3132V|OT9332|NX3132XL}] - Model Number (privilege: advanced)

Selects the switches that match the specified model number.

[-type {cluster-network|management-network}] - Switch Network (privilege: advanced)

Selects the switches that match the specified switch type.

[-sw-version <text>] - Software Version (privilege: advanced)

Selects the switches that match the specified software version.

[-is-monitoring-enabled-operational {true|false}] - Switch Monitoring Status (privilege: advanced)

Selects the switches that match the specified operational monitoring status.

[-reason <text>] - Reason For Not Monitoring (privilege: advanced)

Selects the switches that match the specified reason.

[-version-source <text>] - Source Of Switch Version (privilege: advanced)

Selects the switches that match the specified version source (for example, from SNMP, CDP or ISDP).

[-serial-number <text>] - Serial Number of the Device (privilege: advanced)

Selects the switches that match the specified serial number.

Examples

```
cluster1::> system cluster-switch 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 cluster switches and management switches.

system cluster-switch show

Display the configuration for cluster and management switches

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

Description

The `system cluster-switch show` command displays configuration details for the monitored cluster switches and management switches.

Parameters

{ [-fields <fieldname>,…]

Selects the fields that have the specified name.

| [-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 <text>] - DEPRECATED-Community String or SNMPv3 Username

> This parameter is deprecated and may be removed in a future release of Data ONTAP. Use `-community-or-username` instead.

Selects the switches that match the specified SNMPv2c community string or SNMPv3 username.

[-community-or-username <text>] - Community String or SNMPv3 Username

Selects the switches that match the specified SNMPv2c community string or SNMPv3 username.

[-model {NX5010|NX5020|CAT2960|OTHER|NX5596|CN1610|CN1601|NX3132|OT5548|NX3132V|OT9332|NX3132XL}] - Model Number

Selects the switches that match the specified model number.

[-type {cluster-network|management-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, from SNMP, CDP or ISDP).

[-is-monitoring-enabled-operational {true|false}] - Is Monitored ?

Selects the switches that match the specifed operational monitoring status.

[-serial-number <text>] - Serial Number of the Device

Selects the switches that match the specified serial number.

Examples

```
cluster1::> system cluster-switch 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
cn1601--143-230             management-network 10.238.143.230   CN1601
       Serial Number: 20210200019
        Is Monitored: false
              Reason: Monitoring Disabled by Default
    Software Version: 1.1.0.1
      Version Source: ISDP
cn1601--143-232             management-network 10.238.143.232   CN1601
       Serial Number: 20210200017
        Is Monitored: false
              Reason: Monitoring Disabled by Default
    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 cluster switches and management switches.

```
cluster1::> system cluster-switch 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 cluster switches and management switches.

system cluster-switch polling-interval modify

Modify the polling interval for monitoring cluster and management switch health

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

Description

The `system cluster-switch polling-interval modify` command modifies the interval in which the cluster switch health monitor polls cluster and management 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 cluster-switch polling-interval modify -polling
-interval 41
```

```
Modifies the polling interval of the switches.
```

system cluster-switch polling-interval show

Display the polling interval for monitoring cluster and management switch health

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

Description

The `system cluster-switch polling-interval show` command displays the polling interval used by the health monitor.

Examples

```
cluster1::> system cluster-switch polling-interval show
             Polling Interval (in minutes): 40
```

```
The example above displays the polling interval period for the switches.
```

system cluster-switch threshold show

Display the cluster switch health monitor alert thresholds

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

Description

The `system cluster-switch threshold show` command displays thresholds used by health monitor alerts.

Examples

```
cluster1::> system cluster-switch threshold show
    Per 0.10% values: 1 = 0.10%, 5 = 0.50%
    In Errors Threshold (%) Out Errors Threshold (%)
    ----------------------- ------------------------
                          1                        1
```

Displays the inbound and outbound switch interface packet error thresholds are set at 0.1%. The node platform health monitor also shares the same thresholds in monitoring packet errors of cluster ports on the node.

system configuration commands

system configuration backup copy

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 system configuration backup show command to display configuration backups to copy.

Parameters

`-from-node {<nodename>|local}` - Source Node (privilege: advanced)

Use this parameter to specify the name of the source node where the configuration backup currently exists.

`-backup <text>` - Backup Name (privilege: advanced)

Use this parameter to specify the name of the configuration backup file to copy.

`-to-node {<nodename>|local}` - Destination Node (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

- system configuration backup show

system configuration backup create

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 (privilege: advanced)

Use this parameter to specify the node on which to create the backup file.

`[-backup-name <text>]` - Backup Name (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 (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.
```

system configuration backup delete

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 system configuration backup show command to display saved configuration backups.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the name of the source node where the configuration backup currently exists.

`-backup <text>` - Backup Name (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

- system configuration backup show

system configuration backup download

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 (privilege: advanced)

Use this parameter to specify the name of the node to which the configuration backup is downloaded.

`-source <text>` - Source URL (privilege: advanced)

Use this parameter to specify the source URL of the configuration backup to download.

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

Use this parameter to specify a new local file name for the downloaded configuration backup.

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
```

system configuration backup rename

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 system configuration backup show command to display configuration backups to rename.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the name of the source node where the configuration backup currently exists.

`-backup <text>` - Backup Name (privilege: advanced)

Use this parameter to specify the name of the configuration backup file to rename.

`-new-name <text>` - New Name (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

- system configuration backup show

system configuration backup show

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects configuration backups that are saved on the node you specify.

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

Selects configuration backups that have the backup name you specify.

[-backup-type {node|cluster}] - Backup Type (privilege: advanced)

Selects configuration backups of the type you specify.

[-time <MM/DD HH:MM:SS>] - Backup Creation Time (privilege: advanced)

Selects configuration backups that were saved on the date and time you specify.

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

Selects configuration backups that were saved in the cluster that has the name you specify.

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

Selects configuration backups that have the file size you specify.

[-nodes-in-backup {<nodename>|local}] - Nodes In Backup (privilege: advanced)

Selects configuration backups that include the configuration of the nodes you specify.

[-version <text>] - Software Version (privilege: advanced)

Selects configuration backups that have the software version you specify.

[-is-auto {true|false}] - Backup Created from Schedule (true or false) (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 (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.
```

system configuration backup upload

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.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Use this parameter to specify the name of the node from which the configuration backup is uploaded.

`-backup <text>` - Backup Name (privilege: advanced)

Use this parameter to specify the file name of the configuration backup to upload.

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

Use this parameter to specify the destination URL of 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
```

system configuration backup settings modify

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 (privilege: advanced)

Use this parameter to specify the destination URL for uploads of configuration backups. Use the value "" to remove the destination URL.

`[-username <text>]` - Username for Destination (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 system configuration backup settings set-password command to change the password used with this user name.

`[-numbackups1 <integer>]` - Number of Backups to Keep for Schedule 1 (privilege: advanced)

Use this parameter to specify the number of backups created by backup schedule 1 to keep on the destination system. If the number of backups exceeds this number, the oldest backup is removed.

`[-numbackups2 <integer>]` - Number of Backups to Keep for Schedule 2 (privilege: advanced)

Use this parameter to specify the number of backups created by backup schedule 2 to keep on the destination system. If the number of backups exceeds this number, the oldest backup is removed.

`[-numbackups3 <integer>]` - Number of Backups to Keep for Schedule 3 (privilege: advanced)

Use this parameter to specify the number of backups created by backup schedule 3 to keep on the destination system. If the number of backups exceeds this number, the oldest backup is removed.

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

- system configuration backup settings set-password

system configuration backup settings set-password

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 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 system configuration backup settings show command to display the destination URL for configuration backups. Use the 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

- system configuration backup settings modify
- system configuration backup settings show

system configuration backup settings show

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 system configuration backup settings set-password command to change the password used with the user name on the destination. Use the 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]

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
                                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

- system configuration backup settings set-password
- system configuration backup settings modify

system configuration recovery cluster modify

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 (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
```

system configuration recovery cluster recreate

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 system configuration recovery cluster rejoin command.

Parameters

`-from {node|backup}` - From Node or Backup (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 (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 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

- system configuration recovery cluster rejoin
- system configuration backup show

system configuration recovery cluster rejoin

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 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 (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

- system configuration recovery cluster recreate

system configuration recovery cluster show

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 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 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

- system configuration recovery cluster recreate
- system configuration recovery cluster rejoin

system configuration recovery cluster sync

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 (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.
```

system configuration recovery node restore

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 system configuration backup show command to view available configuration backup files.

Parameters

`-backup <text>` - Backup Name (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 (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 (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

- system configuration backup show

system controller commands

system controller show

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
```

system controller bootmedia show-serial-number

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 name
- 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}] - 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.

Examples

The following example displays the information of the bootmedia devices 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 bootmedia 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
```

system controller bootmedia show

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}] - 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.

Examples

The following example displays the information of the bootmedia devices 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
```

system controller clus-flap-threshold show

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.

system controller config show-errors

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::>
```

system controller config show

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.

[-slot <integer>] - Slot Number

Selects the configuration information that matches the specified slot.

[-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.

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::>
```

system controller config pci show-add-on-devices

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::>
```

system controller config pci show-hierarchy

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::>
```

system controller environment show

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}] - 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
```

system controller flash-cache show

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.

[-slot <integer>] - Slot

If this parameter is specified, only status information for the matching slot is displayed.

[-subslot <integer>] - Subslot

If this parameter is specified, only status information for the matching subslot 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_threshold}] - 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
                Model  Part      Serial               Firmware Capacity
Device
Node       Slot Number Number    Number               Version      (GB)
State
---------- ---- ------ --------- -------------------- -------- --------
------
node1
           6-1  X9172A 119-00209 A22P7061550000004    NA00         4096 ok
           6-2  X9170A 119-00207 A22P5061550000135    NA00         1024 ok
node2
           6-1  X9172A 119-00209 A22P7061550000007    NA00         4096 ok
           6-2  X9170A 119-00207 A22P5061550000091    NA00         1024 ok
4 entries were displayed.
```

system controller flash-cache secure-erase run

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 givien Flash Cache device.

Parameters

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

Selects the node of the specified Flash Cache devices.

`-slot <text>` - Slot

Selects the slot or slot-subslot 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 -slot
0-1
```

system controller flash-cache secure-erase show

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.

`[-slot <text>]` - Slot

If this parameter is specified, only status information for the matching slot is displayed. Slot can have a format of slot or slot-subslot.

`[-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      Slot Device State
--------- ---- ------------
node1
          6-1  ok
          6-2  erasing
node2
          6-1  erased
          6-2  ok
4 entries were displayed.
```

system controller fru show-manufacturing-info

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.

[-fru-subtype <text>] - FRU Sub-type

Selects information about the FRU with the specified FRU subtype.

[-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.
```

system controller fru show

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}] - 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
```

system controller fru led disable-all

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 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

- system controller fru led show

system controller fru led enable-all

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 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

- system controller fru led show

system controller fru led modify

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 system controller fru led show command matches that specified on the FRU map label.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

Selects FRU fault LEDs on the specified nodes.

`-fru-id <FRU LED key>` - FRU ID (privilege: advanced)

Selects the FRU fault LEDs that match the specified FRU type.

`-fru-slot <integer>` - FRU Slot (privilege: advanced)

Selects the FRU fault LEDs that match the specified slot.

`[-fru-state {on|off|unknown}]` - FRU State (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 on
```

The example below turns on all PCI FRU fault LEDs.

```
cluster1::*> system controller led modify -node node1 -fru-id pci -fru
-slot * -fru-state on
```

Related Links

- system controller fru led show

system controller fru led show

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] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects FRU fault LEDs on the specified nodes.

[-fru-id <FRU LED key>] - FRU ID (privilege: advanced)

Selects the FRU fault LEDs that match the specified FRU type.

[-fru-slot <integer>] - FRU Slot (privilege: advanced)

Selects the FRU fault LEDs that match the specified slot.

[-fru-bay <text>] - FRU Bay (privilege: advanced)

Selects the FRU fault LEDs that match the specified bay.

[-fru-state {on|off|unknown}] - FRU State (privilege: advanced)

Selects the FRU fault LEDs that match the specified status.

[-lit-by <text>] - Lit By (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     -
```

system controller ioxm show

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}] - 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
```

system controller location-led modify

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
```

system controller location-led show

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
```

system controller memory dimm show

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

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}] - 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.

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
      Node           Name    Count     Count    Method Socket Channel
Number  Status
      -------------- ------- --------  -------- ------ ------ -------
------  --------
      node1          DIMM-1         0         0 bucket      0       0
0  ok
                     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: ok
                  Unique Name of DIMM: DIMM-1
             Display Name for the DIMM: DIMM-1
        CECC Alert Method: bucket
             Replace DIMM: false
```

system controller nvram-bb-threshold show

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.

system controller pci show

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}] - 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
```

system controller pcicerr threshold modify

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
```

system controller pcicerr threshold show

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
```

system controller platform-capability show

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.
```

system controller service-event delete

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 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

- system controller service-event show

system controller service-event show

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 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 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

- system controller service-event delete

system controller sp config show

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}] - 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
```

system controller sp upgrade show

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}] - 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
```

system feature-usage commands

system feature-usage show-history

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-
000000000000001122334455 -feature-name NFS
Node Serial Number: 1-81-000000000000001122334455
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.
```

system feature-usage show-summary

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: 1-81-0000000000000001122334455
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: 1-81-0000000000000001122334466
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.
```

system fru-check commands

system fru-check show

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.

[-ftype {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.

system ha commands

system ha interconnect config show

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] }

Use this parameter to display all the fields from all nodes in cluster.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display all the fields from the specified node in the cluster.

[-transport <text>] - Interconnect Type (privilege: advanced)

Selects the nodes that match this HA interconnect transport type.

[-local-sysid <integer>] - Local System ID (privilege: advanced)

Selects the nodes that match this local system unique identifier.

[-partner-sysid <integer>] - Partner System ID (privilege: advanced)

Selects the nodes that match this partner system unique identifier.

[-initiator {local|partner}] - Connection Initiator (privilege: advanced)

Selects the nodes that match this parameter value. The value is the initiator of the connection request.

[-port-name <text>,…] - Port (privilege: advanced)

Selects the nodes that match this port name.

[-ipaddress <text>,…] - IP Address (privilege: advanced)

Selects the nodes that match this IP address.

[-interface {backplane|external}] - Interface (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.
```

system ha interconnect link off

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 (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 (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
```

system ha interconnect link on

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 (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 (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
```

system ha interconnect ood clear-error-statistics

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 (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
```

system ha interconnect ood clear-performance-statistics

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 (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
```

system ha interconnect ood disable-optimization

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 (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
```

system ha interconnect ood disable-statistics

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 (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
```

system ha interconnect ood enable-optimization

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 (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
```

system ha interconnect ood enable-statistics

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 (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
```

system ha interconnect ood send-diagnostic-buffer

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 (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
```

system ha interconnect ood status show

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] }

Use this parameter to display all the fields from all nodes in cluster.

[-node {<nodename>|local}] - Node (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 (privilege: advanced)

Selects the nodes that match this parameter value.

[-is-coalescing-enabled {true|false}] - Is Coalescing Enabled (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.
```

system ha interconnect port show

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] }

Use this parameter to display all the fields from all nodes in the cluster.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display all the fields from the specified node in the cluster.

[-link-monitor {on|off}] - Link Monitor Detection (privilege: advanced)

Selects the nodes that match this parameter value.

[-port <integer>,…] - Port Number (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 (privilege: advanced)

Selects the nodes that match this parameter value.

[-link-layer-state {invalid|down|initialize|armed|active|reserved}] - Link Layer State (privilege: advanced)

Selects the nodes that match this parameter value.

[-phy-link-up-count <integer>,…] - Physical Link Up Count (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 (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 (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.
```

system ha interconnect statistics clear-port-symbol-error

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.

To display the high-availability interconnect device port statistics, use the statistics show -object ic_hw_port_stats command.

Parameters

`-node <nodename>` - Node (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

- statistics show

system ha interconnect statistics clear-port

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.

To display the high-availability interconnect device port statistics, use the statistics show -object ic_hw_port_stats command.

Parameters

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

Selects the nodes that match this parameter value.

Examples

```
cluster1::*> system ha interconnect statistics clear-port -node ic-f8040-
01
```

Related Links

- statistics show

system ha interconnect statistics show-scatter-gather-list

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] }

Use this parameter to display all the fields from all nodes in cluster.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display all the fields from the specified node in the cluster.

[-sge <integer>,…] - Scatter-Gather Entry (privilege: advanced)

Selects the nodes that match this scatter-gather element index value.

[-total-count <integer>,…] - Total Count (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 (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.
```

system ha interconnect statistics performance show

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] }

Use this parameter to display all the fields from all nodes in cluster.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display all the fields from the specified node in the cluster.

[-elapsed <integer>] - Elapsed Time (secs) (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 (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) (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 (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 (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 (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 (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) (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 (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 (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) (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 (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) (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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) (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 (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 (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 (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) (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) (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 (privilege: advanced)

Selects the nodes that match this parameter value. The value is the maxmimum length of the scatter-gather list supported by the interconnect device.

[-ic-recq-waits <integer>] - Total Receive Queue Waits to Post Buffer (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) (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.
```

system ha interconnect status show

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] }

Use this parameter to display all the fields for the specified node or all the nodes.

[-node {<nodename>|local}] - Node (privilege: advanced)

Use this parameter to display all the fields for the specified node.

[-link-status {up|down}] - Link Status (privilege: advanced)

Selects the nodes that match this parameter value. The value *up* means link is online.

[-link0-status {up|down}] - Link 0 Status (privilege: advanced)

Selects the nodes that match this parameter value. The value *up* means link is online.

[-link1-status {up|down}] - Link 1 Status (privilege: advanced)

Selects the nodes that match this parameter value. The value *up* means link is online.

[-ic-rdma {up|down}] - IC RDMA Connection (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 (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 (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 (privilege: advanced)

Selects the nodes that match this PCI slot number.

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

Selects the nodes that match this interconnect device driver name.

[-firmware <text>] - Firmware Revision (privilege: advanced)

Selects the nodes that match this firmware version.

[-version <text>] - Version Number (privilege: advanced)

Selects the nodes that match this parameter value.

[-device-type <text>] - Device Type (privilege: advanced)

Selects the nodes that match this interconnect device type.

[-serial-number <text>] - Serial Number (privilege: advanced)

Selects the nodes that match this interconnect device serial number.

[-debug-firmware {yes|no}] - Debug Firmware (privilege: advanced)

Selects the nodes that match this parameter value.

[-command-revision <integer>] - Command Revision (privilege: advanced)

Selects the nodes that match this interconnect device command revision.

[-hardware-revision <integer>] - Hardware Revision (privilege: advanced)

Selects the nodes that match this interconnect device hardware revision.

[-port1 <integer>] - Port Number 1 (privilege: advanced)

Selects the nodes that match this parameter value.

[-port1-port-name <text>] - Port Name (privilege: advanced)

Selects the nodes that match this port name.

[-port1-gid <text>] - Global Identifier (privilege: advanced)

Selects the nodes that match this global identifier value.

[-port1-base-lid <text>] - Base Local Identifier (privilege: advanced)

Selects the nodes that match this base local identifier value.

[-port1-rm-lid <text>] - Remote Local Identifier (privilege: advanced)

Selects the nodes that match this remote local identifier value.

[-port1-mtu <integer>] - Maximum Transmission Unit (privilege: advanced)

Selects the nodes that match this parameter value.

[-port1-data-rate <text>] - Data Rate (privilege: advanced)

Selects the nodes that match this parameter value.

[-port1-link-info <text>] - Link Information (privilege: advanced)

Selects the nodes that match this parameter value.

[-port1-qsfp-vendor <text>] - QSFP Vendor (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) vendor name.

[-port1-qsfp-part-number <text>] - QSFP Part Number (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) part-number.

[-port1-qsfp-type <text>] - QSFP Type (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) type.

[-port1-qsfp-serial-number <text>] - QSFP Serial Number (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) serial number.

[-port2 <integer>] - Port Number 2 (privilege: advanced)

Selects the nodes that match this parameter value.

[-port2-port-name <text>] - Port Name (privilege: advanced)

Selects the nodes that match this port name.

[-port2-gid <text>] - Global Identifier (privilege: advanced)

Selects the nodes that match this global identifier value.

[-port2-base-lid <text>] - Base Local Identifier (privilege: advanced)

Selects the nodes that match this base local identifier value.

[-port2-rm-lid <text>] - Remote Local Identifier (privilege: advanced)

Selects the nodes that match this remote local identifier value.

[-port2-mtu <integer>] - Maximum Transmission Unit (privilege: advanced)

Selects the nodes that match this parameter value.

[-port2-data-rate <text>] - Data Rate (privilege: advanced)

Selects the nodes that match this parameter value.

[-port2-link-info <text>] - Link Information (privilege: advanced)

Selects the nodes that match this parameter value.

[-port2-qsfp-vendor <text>] - QSFP Vendor (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) vendor name.

[-port2-qsfp-part-number <text>] - QSFP Part Number (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) part number.

[-port2-qsfp-type <text>] - QSFP Type (privilege: advanced)

Selects the nodes that match this QSFP (Quad Small Form-factor Pluggable) type.

[-port2-qsfp-serial-number <text>] - QSFP Serial Number (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
```

system health commands

system health alert delete

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 *
```

system health alert modify

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
```

system health alert show

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
           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
```

system health alert definition show

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
```

system health autosupport trigger history show

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
```

system health config show

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
```

system health policy definition modify

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 *
```

system health policy definition show

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
```

system health status show

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
```

system health subsystem show

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 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>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

- system health alert show

system license commands

system license add

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.

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
```

system license clean-up

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.

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:

```
cluster1::> system license clean-up -n -unused -expired
Serial number: 1-80-000011
Owner: cluster1
Package                   Reason
------------------------
-------------------------------------------------------
SnapLock                  Demo License has expired
SnapProtectApps           Demo License has expired

Serial number: 1-81-0000000000000004062522917
Owner: none
Package                   Reason
------------------------
-------------------------------------------------------
NFS                       Serial number is not used by any node in the
cluster
CIFS                      Serial number is not used by any node in the
cluster
```

The following example deletes the expired licenses:

```
cluster1::> system license clean-up -expired
2 demo licenses deleted.
```

The following example deletes the unused licenses:

```
cluster1::> system license clean-up -unused
2 unused licenses deleted.
```

system license delete

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 1-81-000000000000000000000123456 from the cluster:

```
cluster1::> system license delete -serial-number 1-81-
000000000000000000000123456 -package CIFS
```

system license show-status

Display license status

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

Description

This command displays the status of all Data 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}] - 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}] - 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.
```

system license show

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}] - Type

If you use this parameter, the command displays information only about the licenses that have the license type you specify.

[-legacy {yes|no}] - Legacy

If you use this parameter, the command displays information only about the licenses that matches the legacy field 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.

Examples

The following example displays default information about all licensed packages in the cluster:

```
cluster1::> system license show
Serial Number: 1-80-123456
Owner: cluster1
Package           Type    Description           Expiration
----------------- ------- --------------------- --------------------
Base              site    Cluster Base License  -
NFS               site    NFS License           -
iSCSI             site    iSCSI License         -

Serial Number: 1-81-00000000000000001122334455
Owner: node1
Package           Type    Description           Expiration
----------------- ------- --------------------- --------------------
NFS               license NFS License           -
SnapRestore       license SnapRestore License   -
5 entries were displayed.
```

system license capacity show

(DEPRECATED)-Show license capacity status

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 Data ONTAP. Use the "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: 1-81-00000000000012345678901234567
                              Max  Current
Package                  Capacity Capacity Expiration
------------------------ -------- -------- -------------------
Select                        2TB  15.81GB 4/11/2016 00:00:00
Node:          node2
Serial Number: 1-81-00000000000000000000123456788
                              Max  Current
Package                  Capacity Capacity Expiration
------------------------ -------- -------- -------------------
Select                          -  10.40TB 4/11/2016 00:00:00
Node:          node3
Serial Number: 1-81-00000000000000000000123456789
                              Max  Current
Package                  Capacity Capacity Expiration
------------------------ -------- -------- -------------------
Select                          -  10.40TB 4/11/2016 00:00:00
Node:          node4
Serial Number: 1-81-00000000000012345678901234567
                              Max  Current
Package                  Capacity Capacity Expiration
------------------------ -------- -------- -------------------
Select                        2TB  15.81GB 4/11/2016 00:00:00
```

Related Links

- system license show-status

system license entitlement-risk show

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 <Node Serial Number>] - 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
```

system license status show

(DEPRECATED)-Display license status

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 Data ONTAP. Use the "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}] - 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

- system license show-status

system node commands

system node halt

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. In a two-node MetroCluster configuration, this parameter prevents automatic unplanned switchover.

> If -inhibit-takeover is set to true, the default behavior as seen with command storage failover show` -fields onreboot` 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.

Examples

The following example shuts down the node named cluster1 for hardware maintenance:

```
cluster1::> system halt -node cluster1 -reason 'hardware maintenance'
```

Related Links

- storage failover show

system node migrate-root

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 (privilege: advanced)

Specifies the node that owns the root aggregate that you wish to migrate. The value *`local`* specifies the current node.

{ `-disklist <disk path name>,…` - List of Disks for New Root Aggregate (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 Type for the New Root Aggregate (privilege: advanced)

Specifies the RAID type of the root aggregate. The default value is *`raid-dp`* .

| `-resume <true>` - Resume a Failed Migrate Operation (privilege: advanced) }

Resumes a failed migrate-root operation if the new_root aggregate is created and the old root aggregate is in the retricted 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
```

system node modify

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 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 (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 (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-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.

[-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

- system node show

system node reboot

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. 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.

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"
```

system node rename

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 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

- system node show

system node restore-backup

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 Data 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 (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 (privilege: advanced)

Specifies the IP address for the target node.

system node revert-to

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 (privilege: advanced)

Specifies the node that is to be reverted. The value `local` specifies the current node.

`-version <revert version>` - Data ONTAP Version (privilege: advanced)

Specifies the version of Data ONTAP to which the node is to be reverted.

`[-check-only <true>]` - Capability Check (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 Data ONTAP version 9.2

```
cluster1::*> system node revert-to -node node1 -version 9.2
```

system node run-console

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 Data 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::>
```

system node run

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.

> 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
```

system node show-discovered

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 wish 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 Data 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
```

system node show-memory-errors

Display Memory Errors on DIMMs

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

Description

`system node show-memory-errors` prints the history of memory (storage controller's RAM) errors since boot. This command can be useful in diagnosing memory problems or determining which DIMM, if any, might need replacement. Some correctable ECC errors are to be expected under normal operation, but many occurring on a particular DIMM might indicate a problem. All the fields are read only and can be used to filter the output. The maximum number of physical address and timestamps reported is 160.

Parameters

{ [-fields <fieldname>,…]

If you specify 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)

The `-verbose` parameter enables verbose mode, resulting in the display of more detailed output.

| [-instance] }

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

[-node {<nodename>|local}] - Node (privilege: advanced)

When provided, the `-node` parameter specifies the nodes for which the memory error statistics are to be displayed. When the `-node` is not provided, the command is applied to all the nodes in the cluster.

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

This parameter refers to the DIMM ID. It can be used to look at the correctable ECC error count on a specific DIMM.

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

This parameter specifies the DIMM name for which the memory error statistics are to be displayed.

[-cecc <integer>] - Correctable ECC Error Count (privilege: advanced)

This parameter can be used to get all the DIMMs with the specified correctable ECC error count.

[-merr {true|false}] - Multiple Errors on Same Address (privilege: advanced)

Use this parameter with the values true to specify whether the error was seen multiple times on the same physical address. It can also be used to look at all the DIMMs with multiple errors on same address.

[-timestamp <text>,…] - Error Time (privilege: advanced)

This specifies the time at which the error was seen on the DIMM.

[-addr <text>,…] - Error Address (privilege: advanced)

This specifies the physical address on which the error was seen.

Examples

```
cluster1::*> system node show-memory-errors
  Correctable ECC Memory Errors:
Node: localhost
DIMM      CECC  Multiple Err
  Name     Count  Same Address
  ------- ------  ------------
  DIMM-1       0  false
  DIMM-2       0  false
  DIMM-3       0  false
  DIMM-4       0  false
  DIMM-5       4  true
  DIMM-6       1  false
  DIMM-7       1  false
  DIMM-8       0  false
  8 entries were displayed.
cluster1::*> system node show-memory-errors -verbose
  Correctable ECC Memory Errors:
Node: localhost
DIMM      CECC  Multiple Err                              Physical
  Name     Count  Same Address  Timestamp                   Address
  ------- ------  ------------  --------------------  -------------
  DIMM-1       0  false                             -              -
  DIMM-2       0  false                             -              -
  DIMM-3       0  false                             -              -
  DIMM-4       0  false                             -              -
  DIMM-5       4  true           12/02/2013 08:17:43         0xD640
                                 12/02/2013 08:17:57     0x3F7FF800
                                 12/02/2013 08:18:03    0x11743D000
                                 12/02/2013 08:18:37    0x11743D000
  DIMM-6       1  false          12/02/2013 08:17:53        0x87EC0
  DIMM-7       1  false          12/02/2013 08:17:51    0x13DED8900
  DIMM-8       0  false                      -              -
  8 entries were displayed.
```

system node show

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 wish to view. Use the 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 (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 (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.

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

- system node modify

system node autosupport invoke-core-upload

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 "http" and "https" protocols are supported. If this parameter is omitted, the message is sent to all 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 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 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

- system node autosupport modify
- system node coredump show

system node autosupport invoke-performance-archive

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.

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," "http," 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
```

system node autosupport invoke-splog

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 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 Procesor 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," "http," 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.

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

- system node autosupport modify

system node autosupport invoke

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 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", "http", 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 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

- system node autosupport modify

system node autosupport modify

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 a node.

Parameters

-node {<nodename>|local} - Node

Use this parameter to specify the node being configured.

[-state {enable|disable}] - State

Use this parameter to specify whether AutoSupport is enabled or disabled on the node. 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 the node sends AutoSupport messages. 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 with no 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 *25* .

Also, you can optionally prepend a user name and password combination for authentication to each mail server. The format of the username and password pair is *user1:pw1@mymailhost.example.com* . The username and password can be specified on none, all, or some of the mail hosts.

The default value for this parameter is *mailhost* .

[-from <mail address>] - From Address

Use this parameter to specify the e-mail address from which the node sends AutoSupport messages. The default is Postmaster@xxx where xxx is the name of the system.

[-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 system node autosupport trigger modify command. By default, no list is defined.

[-noteto <mail address>,…] - (DEPRECATED) List of Noteto Addresses

This parameter has been deprecated and might be removed in a future version of Data ONTAP.

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 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|http|https}] - Protocol to Contact Support

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 http or 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 or HTTPS proxy if the `-transport` parameter is set to HTTP or HTTPS and your organization uses a proxy. Enter the URL without an http:// or https:// prefix. If authentication is required, use the format "[username]@[host][:[port]]". You will be prompted for the password. 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` .

[-nht {true|false}] - NHT Enable

Use this parameter to specify whether NHT disk drive health data is sent to technical support and addresses specified in the `-partner-address` parameter when disk drives fail. The default 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 HTTP 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 HTTP and HTTPS transfers. This parameter applies only to messages sent to technical support and only if the -transport parameter is set to HTTP or 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.

Eliminated data might 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, and qtree names.

The default is `false` .

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

Use this parameter with the value `true` to force the 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 security config ocsp enable -app autosupport and security config ocsp disable -app autosupport .

[-ondemand-state {enable|disable}] - AutoSupport OnDemand State (privilege: advanced)

Use this parameter to specify whether the AutoSupport OnDemand feature is enabled or disabled on the node. 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 (privilege: advanced)

Use this parameter to specify whether the AutoSupport OnDemand Remote Diagnostics feature is enabled or disabled on the node. 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.

Examples

The following example enables AutoSupport on a node named node3 with the following settings:

- SMTP mail host named smtp.example.com.
- E-mail "from" address of alerts@node3.example.com
- E-mail "to" address of support@example.com
- AutoSupport messages sent to support personnel
- HTTPS set as transport protocol
- Short-note address of pda@example.com
- If sending fails, the system will wait 23 minutes before retrying.

```
cluster1::> system node autosupport modify -node node3 -state enable -mail
-hosts smtp.example.com -from alerts@node3.example.com -to
support@example.com -support enable -transport https -noteto
pda@example.com -retry-interval 23m
```

The following examples show how to modify AutoSupport URLs when using IPv6 address literals:

```
cluster1::> system node autosupport modify -node node1 -mail-hosts
[2620:10a:4002:6004::bbbb]:25
```

```
cluster1::> system node autosupport modify -node node1 -proxy-url
username:password@[2620:10a:4002:6004::bbbb]:8080
```

Related Links

- system node autosupport trigger modify
- security config ocsp enable
- security config ocsp disable

system node autosupport show

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.

| [-nht-performance]

Use this parameter to display NHT and performance information about 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 for each node in the cluster, and identify the transport protocol and the support proxy URL used by each node.

| [-support-smtp]

Use this parameter to display whether SMTP (e-mail) support is enabled for each node in the cluster, and identify the transport protocol and the "to" mail address used by each node.

| [-instance] }

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

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

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.

[-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

This parameter has been deprecated and might be removed in a future version of Data 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|http|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 HTTP/HTTPS

Use this parameter to display information only about nodes that use the URL you specify to send messages through HTTP and HTTPS POST operations.

[-put-url <text>] - Support URL for HTTP/S PUT

Use this parameter to display information only about nodes that use the URL you specify to send messages through HTTP 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.

[-nht {true|false}] - NHT Enable

Use this parameter with the value "true" to display information only about nodes that send NHT disk drive data. Use this parameter with the value "false" to display information only about nodes that do not send NHT data.

[-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>] - Last Subject Sent

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>] - Last Time Sent

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 HTTP 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

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 (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 (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.

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
                                 From Address: alerts@node3.example.com
                            List of To Addresses: support@example.com
                       List of Noteto Addresses: -
                     List of Partner Addresses: partner@node4.example.com
Send AutoSupport Messages to Vendor Support: enable
                  Protocol to Contact Support: https
                            Support Proxy URL: support.proxy.example.com
                             Hostname Subject: true
                                   NHT Enable: 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: 3/11/2011 06:00:03
                             Maximum HTTP Size: 50MB
                             Maximum SMTP Size: 5MB
                          Remove Sensitive Data: false
      Validate Digital Certificate Received: true
   Continue Local Collection while Disabled: true
```

system node autosupport check show-details

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: http-https
          Component: http-put-destination
             Status: ok
             Detail: Successfully connected to "support.netapp.com/put".
Component: http-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.
```

system node autosupport check show

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 HTTP/HTTPS Destinations

Selects the nodes that match this parameter value. This parameter specifies whether connectivity to the AutoSupport HTTP 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
```

system node autosupport destinations show

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 HTTP POST, HTTP 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
```

system node autosupport history cancel

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 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

- system node autosupport history retransmit

system node autosupport history retransmit

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 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 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

- system node autosupport history show

system node autosupport history show-upload-details

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

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>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.
```

system node autosupport history show

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

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>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.
```

system node autosupport manifest show

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.
- `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.
```

system node autosupport trigger modify

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

This parameter has been deprecated and might be removed in a future version of Data ONTAP.

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
```

system node autosupport trigger show

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 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

This parameter has been deprecated and might be removed in a future version of Data 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

- system node autosupport destinations show

system node coredump delete-all

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
```

system node coredump delete

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
```

system node coredump save-all

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
```

system node coredump save

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 system node coredump config modify . A saved core dump can be uploaded to a remote site for support analysis; see the 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

- system node coredump config modify
- system node coredump upload

system node coredump show

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.
```

system node coredump status

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
```

system node coredump trigger

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` , system node halt , or 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 Data ONTAP running on the nodes. This command is helpful in cases where Data 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 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 (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 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

- system node halt
- system node reboot
- system node coredump show

system node coredump upload

(DEPRECATED)-Upload a coredump to a remote site

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 of Data ONTAP. Use "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 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

- system node autosupport invoke-core-upload
- system node coredump config modify

system node coredump config modify

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

This option is deprecated and might be removed in a future release of Data ONTAP. Use the `-uri` parameter of the "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

- system node autosupport invoke-core-upload

system node coredump config show

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

This option is deprecated and might be removed in a future release of Data ONTAP. Use the `-uri` parameter of the "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

- system node autosupport invoke-core-upload

system node coredump reports delete

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
```

system node coredump reports show

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 thenode 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
```

system node coredump reports upload

(DEPRECATED)-Upload an application core report to a remote site

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 of Data 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
```

system node coredump segment delete-all

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
```

system node coredump segment delete

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
```

system node coredump segment show

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
```

system node environment sensors show

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}] - 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}] - 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
[...]
```

system node external-cache modify

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` .

system node external-cache show

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.

system node firmware download

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 (privilege: advanced)

This specifies the node or nodes on which the firmware is to be updated.

-package <text> - Package URL (privilege: advanced)

This parameter specifies the URL that provides the location of the package to be fetched. Standard URL schemes, including HTTP, 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.

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
```

system node hardware nvram-encryption modify

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.

> 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
```

system node hardware nvram-encryption show

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.
```

system node hardware tape drive show

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.
```

system node hardware tape library show

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.
```

system node hardware unified-connect modify

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 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.

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

- network fcp adapter modify

system node hardware unified-connect show

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.
```

system node image abort-operation

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 (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
```

system node image get

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 (privilege: advanced)

This parameter specifies the node that fetches and stores the package.

-package <text> - Package URL (privilege: advanced)

This parameter specifies the URL that provides the location of the package to be fetched. Standard URL schemes, including HTTP, 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.

[-replace-package <true>] - Replace the Local File (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 (privilege: advanced)

Use this parameter to enter a package name that is different than the file name in the URL.

[-background <true>] - Run in the background (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
```

system node image modify

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 (privilege: advanced)

This specifies the node on which the software image is located.

`-image {image1|image2|remote}` - Image Name (privilege: advanced)

This specifies the software image that is to be set or unset as the default.

`[-isdefault {true|false}]` - Is Default Image (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.
```

system node image show-update-progress

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 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 (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 (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

- system node image update

system node image show

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] }

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

[-node {<nodename>|local}] - Node (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 (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 (privilege: advanced)

Selects information about the software images with the specified default setting.

[-iscurrent {true|false}] - Is Current Image (privilege: advanced)

Selects information about the software images that have the specified currency value.

[-kernel-path <text>] - Kernel Image Path (privilege: advanced)

Selects information about the software images that have the specified kernel image path.

[-rootfs-path <text>] - Root Filesystem Image Path (privilege: advanced)

Selects information about the software images that have the specified root file system image path.

[-version <text>] - Software Version (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 (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.
```

system node image update

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 (privilege: advanced)

This specifies the node on which the software image is located.

`-package <text>` - Package URL (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.

> 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 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 system node image get . For example, *`image.tgz`* . Available packages can be displayed using 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 (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 (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 (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 (privilege: advanced)

Use this parameter to enter a package name that is different than the file name in the URL.

[-background <true>] - Run in the Background (privilege: advanced)

This parameter will allow the operation to run in the background. The progress of the operation can be checked with the command 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 (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

- system node image get
- system node image package show
- system node image show-update-progress

system node image package delete

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 (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 (privilege: advanced)

This parameter specifies the package to be deleted.

Examples

```
::> system image package delete image.tgz
1 entry was deleted.
```

system node image package show

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] }

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

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

system node image package external-device delete

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 (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 (privilege: advanced)

This parameter specifies the file to be deleted.

-device <usb0> - Device (privilege: advanced)

This parameter specifies the name of the external device. Currently, only usb0 is supported.

Examples

```
::> system image package external delete -package image.tgz
```

system node image package external-device show

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] }

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

[-node {<nodename>|local}] - Node (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 (privilege: advanced)

This parameter specifies the file for which the information is displayed.

[-device <usb0>] - Device (privilege: advanced)

This parameter specifies the name of the external device. Currently, only usb0 is supported.

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.
```

system node internal-switch show

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.
```

system node internal-switch dump eeprom

Display onboard switch eeprom config

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

Description

The `system node internal-switch dump eeprom` command is used to show the internal switch eeprom firmware content.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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 internal switch resides on.

[-switch-id <integer>] - Switch

Use this parameter to specify the switch id. For example, 1.

[-version <text>] - Version

Use this parameter to specify the version of firmware needs to be dumped.

[-eeprom <text>] - Firmware

Use this parameter to specify the firmware content needs to be dumped.

Examples

The following example shows the internal switch eeprom firmware content of the internal switch 0 on Node1.

```
cluster1::> system node internal-switch dump eeprom -node Node1 -switch-id
0
EEPROM Configuration
 --------------------
Node: Node1    , Switch: 0 , Version: Version2 SB_XX Switch0
7fa300007fa900007fa800007fa700007fa610007fa5d0007fa200007fa30000
7fa610107fa802337fa713137fa5b0007fa200007fa300007fa610207fa80233
7fa731317fa5b0007fa200007fa300007fa610407fa802117fa733337fa5b000
8027102080282c808047101080482c808087102080882c808107101081082c80
8207104082082c808407104084082c8088082c808809000a880a80fa8024007c
7fe990007fea00047fe9bf017fea00047fe9b0007fea00197fe9b0017fea0020
7fe9b0027fea4c887fe9b0037fea01197fe9b0047feafff07fe9b0057fea00ff
7fe9b0067fea10007fe9b0078024007fffff56657273696f6e20320000005342
5f5858000005377697463682030000fffffff0040e0e8100065b4100065b4
100065b42b2adfe000405860000000000000012b2adfe0100087700040e0e8
000003b6000400087fc5e9a0000000010008770000000000000003b600040008
10008770100014682b2adfe00000000010001050000000017fc5e9a010001468
2ab094700040e0e8100065b4100065b42b2adfe0004059bc000000c0100018e8
7fc5e9d80040c01800000000000000c0000000260000000010008770000000000
000003b6000e001c10008770000000010000001e0000001e2ab094700040e0e8
7fc5ea18004049100000000000000000060041bc2c7fc5
```

system node internal-switch dump port-mapping

Display onboard switch port mapping

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

Description

The `system node internal-switch dump port-mapping` command is used to show the internal switch port connections.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields <fieldname>, …` parameter, the command output also includes the specified field or fields. You 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.

[-role <text>] - Port Name

Use this parameter to specify the port name. For example, e0M.

Examples

The following example shows the port connections of the internal switch 0 on Node1.

```
cluster1::> system node internal-switch dump port-mapping -node Node1
-switch-id 0
Port  Port Name
----  -----------------
Node: Node1    , Switch: 0
0     sw-wrench
1     sw-locked-wrench
2     sw-e0M
3     sw-e0P
4     sw-midplane-1
5     sw-expander-1
6     sw-sp-1
7 entries were displayed.
```

system node internal-switch dump stat

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
```

system node power on

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 (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::*>
```

system node power show

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::>
```

system node root-mount create

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 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

- system node root-mount show

system node root-mount delete

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 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

- system node root-mount show

system node root-mount show

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.
```

system node upgrade-revert show

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. 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] }

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 status information only about upgrades or reversions that are slated to occur on the nodes you specify.

[-upgrade-version <integer>] - Cluster Upgrade Version (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 (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 (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 (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 (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 (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 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 cluster show and cluster ping-cluster to ensure that all nodes are healthy and in communication.

[-node-status-msg <text>] - Node Status Message (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

- system node upgrade-revert upgrade
- cluster show
- cluster ping-cluster

system node upgrade-revert upgrade

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 (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.
```

system node virtual-machine hypervisor modify-credentials

Modify hypervisor IP address and its credentials

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

Description

The `system node virtual-machine hypervisor modify-credentials` command is used to set the *`IP address`* of the hypervisor on which the node is running or vSphere credentials (i.e. `-new-username` or `-new-password`).

Parameters

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

Name of the Data ONTAP node running in a virtual machine. It is a required field and the node must exist in the cluster.

`[-new-server <text>]` - New Hypervisor IP Address

New vSphere server controlling this virtual machine. It can be either an IP address or (if name resolution is enabled) a hostname.

`[-new-username <text>]` - New Hypervisor Username

New vSphere username for the `-new-server` specified above.

`[-new-password <text>]` - New Hypervisor Password

New vSphere password for the `-new-server` specified above.

Examples

The following example sets the IP address and the credentials of the vSphere server on which the node is running.

```
cluster1::> system node virtual-machine hypervisor modify-credentials
 -node node1 -new-server 192.168.0.1 -new-username admin -new-password
pass123
```

system node virtual-machine hypervisor show-credentials

Display hypervisor IP address and its credentials

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

Description

The `system node virtual-machine hypervisor show-credentials` command displays the current vSphere authentication information (except the password). This consists of the vSphere server and username. The vSphere password is not displayed for security reasons. vSphere authentication information is required to use (system node virtual-machine hypervisor show`-node` *<node>* , `system node virtual-machine instance show-system-disks`-node*<node>* and `storage disk show -virtual-machine-disk -info`-node*<node>*) to be able to gather information about the physical host machine. It also attempts to verify the current vSphere authentication information with the vSphere host. If the check succeeds and the credentials are correct, the command displays the following information. If you want to see details about a single node, use the `-node` parameter.

- Node
- Hypervisor name or IP Address
- vSphere Username
- Credentials Correct?: *true*

If the check fails or credentials are incorrect, the command displays an additional *Error* .

- Node
- Hypervisor name or IP Address
- vSphere Username
- Credentials Correct?: *false*
- 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] }

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 Data 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.

[-server <text>] - Hypervisor IP Address or Hostname

Use this parameter to only display the Data ONTAP nodes in the cluster whose vSphere server matches this value.

[-username <text>] - Hypervisor Username

Use this parameter to only display the Data ONTAP nodes in the cluster whose vSphere username matches this value.

[-are-credentials-correct {true|false}] - Credentials Correct?

Get a list of Data ONTAP nodes running with either incorrect (false) or correct (true) vSphere credentials.

[-error <text>] - Error

Get a list of nodes with the specified error. This parameter is most useful when entered with wildcards.

Examples

The following example shows the vSphere server and vSphere username. It also displays whether the server address or its credentials are correct and displays an error if they are not.

```
cluster1::> system node virtual-machine hypervisor show-credentials
                  Node: node1
Hypervisor IP Address: 192.168.0.1
     vSphere Username: administrator
 Credentials Correct?: true
Node: node2
Hypervisor IP Address:
     vSphere Username: admin
 Credentials Correct?: false
Error: [13166] could not find IP addr for host.
Correct the vSphere credentials with the
"system node virtual-machine hypervisor modify-credentials -node" command.

2 entries were displayed.
```

Related Links

- system node virtual-machine hypervisor show
- system node virtual-machine instance show-system-disks
- storage disk show

system node virtual-machine hypervisor show

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 Data 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 Data 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 Data 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 Data 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 is comprised 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 Data 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(R) Xeon(R) 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.
```

system node virtual-machine instance show-system-disks

Display information about Data ONTAP-v 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 the system disks information of the Data ONTAP node running on the virtual machine. There are three types of Data ONTAP-v system disks - boot, core and log. The details contain the physical properties of the disk and its backing information (e.g. backing store name, iSCSI Lun UUID, etc). 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 Name

The name of the Data 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.

[-vmdisk-type <text>] - Type of the System Disk

The type of the system disk, for example "Boot", "Core" or "Log". This optional parameter selects the information about the system disk with the specified type.

[-vmdisk-name <text>] - Data ONTAP Supplied Name of the System Disk

The Data ONTAP-supplied name of the system disk. This optional parameter selects the information about the system disk with the specified name.

[-vmdisk-serial-number <text>] - Data ONTAP Supplied Serial Number of the System Disk

The Data ONTAP-supplied serial number of the system disk. This optional parameter selects the information about system disk with the specified serial number.

[-vmdisk-capacity {<integer>[KB|MB|GB|TB|PB]}] - Size of the System Disk

The size of the system disk. This optional parameter selects the information about system disks with the specified size.

[-vmdisk-area-name <text>] - Name of the System Disk Backing Store

The 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. This optional parameter selects the information about the system disks that reside on the specified backing store.

[-vmdisk-file-name <text>] - File Name of the System Disk Used By the Hypervisor

The file name of the virtual disk used by the hypervisor. Each Data ONTAP disk is mapped to a unique VM disk file. This optional parameter selects the information about the system disk mapped to the specified file name.

[-vmdisk-backing-store-type <text>] - Type of the System Disk Backing Store

The type of the disk backing store. It can be a VMFS volume, a directory on network-attached storage, or a local file system path. This optional parameter selects the information about the system disks that are backed by the specified type.

[-vmdisk-backing-store-capacity {<integer>[KB|MB|GB|TB|PB]}] - Size of the System Disk Backing Store

The size of the disk backing store. This optional parameter selects the information about the system disks that are backed by stores with the specified capacity.

[-vmdisk-backing-store-nas-path <text>] - Full Path to the Backing Store for NAS

The full path to the backing store for network-attached storage. This optional parameter selects the information about the system disks with the specified NAS path.

[-vmdisk-backing-adapter-pci <text>] - Backing Adapter PCI Device ID

The backing adapter PCI device ID for the virtual disk. This optional parameter selects the information about the system disks that have the specified ID as their backing adapter PCI ID.

[-vmdisk-backing-adapter-device <text>] - Name of the Backing Adapter Device

The name of the backing adapter device. This optional parameter selects the information about the system disks that have the specified name as their backing adapter device name.

[-vmdisk-backing-adapter-model <text>] - Type of the Backing Adapter Model

The type of the backing adapter model. This optional parameter selects the information about the system disks that have the specified type as their backing adapter model type.

[-vmdisk-backing-adapter-driver <text>] - Backing Adapter Driver Name of the Initiator

The backing adapter driver name of the initiator. This optional parameter selects the information about the system disks that have the specified name as their initiator’s backing adapter driver name.

[-vmdisk-backing-target-iscsi-name <text>] - iSCSI Name of the System Disk Backing Target

The iSCSI name of the disk backing target. This field is valid only for iSCSI connections. This optional parameter selects the information about the system disks that have the specified name as their backing device iSCSI target name.

[-vmdisk-backing-target-iscsi-address <text>] - iSCSI IP Address of the System Disk Backing Target

The iSCSI IP address of the disk backing target. This field is valid only for iSCSI connections. This optional parameter selects the information about the system disks that have the specified IP as their backing device iSCSI target IP address.

[-vmdisk-backing-device-address <text>] - SCSI Device Name (target-id:lun-id) for the Backing Disk

The SCSI device name for the backing disk. It takes the form target-id:lun-id, for example "2:1". This optional parameter selects the information about the system disks that have the specified backing device address.

[-vmdisk-backing-device-uuid <text>] - Hypervisor-Assigned Unique ID of the Backing Device

The hypervisor-assigned unique ID of the backing device (disk or LUN). This optional parameter selects the information about system disks that are backed by a device with the specified UUID.

[-vmdisk-backing-device-partition <integer>] - Backing Disk Partition Number for the VM Disk File

The backing disk partition number where the corresponding VM disk file resides. This optional parameter selects the information about the system disks that reside on the specified partition number on any disk.

[-vmdisk-backing-device-capacity {<integer>[KB|MB|GB|TB|PB]}] - Size of the Backing Device (Disk or LUN)

The size of the backing device (disk or LUN) in bytes. This optional parameter selects the information about the system disks that have backing devices with the specified capacity.

Examples

The following example shows typical output from the `system node virtual-machine instance show-system-disks` command for the Data ONTAP node running on a virtual machine.

```
cluster1::> system node virtual-machine instance show-system-disks
     Disk Disk  Disk     Store
Node Type Name  Capacity Name   VM Disk File Name
---- ---- ----- -------- ------ ------------------------------
node1
     Boot ad0   1.032 GB store1 store1/node1/DataONTAP.vmdk
     Core ad1   1.505 GB store2 store2/node1/DataONTAP-var.vmdk
     Log  ad2   5.001 GB store3 store3/node1/DataONTAP-nvram.vmdk

Warning: Unable to list entries on node node2. Access to the
         vSphere server is failing since the server hostname or IP
         address is not set. Correct the vSphere credentials with the
         "system node virtual-machine hypervisor modify-credentials"
         command.
3 entries were displayed.
```

system node virtual-machine instance show

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 Data 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 Data 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

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

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

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 Data ONTAP in the cloud.

[-instance-type <text>] - Specifies System Attributes and Use Cost

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

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>] - Version of This VM Instance Information

The version of the Instance Metadata or Agent Wire Protocol as defined by the cloud provider.

[-availability-zone <text>] - Distinct Location within a Region

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

The primary IP address assigned to this virtual machine instance.

[-deployment-id <text>] - Deployment ID of This Instance

A cloud provider-supplied unique deployment ID for this virtual machine, for example "2831c724-97ca-4395-b8d3-a65c2a65b502".

[-fault-domain <integer>] - Fault Domain of This Instance

A cloud provider-assigned numerical fault domain ID for this virtual machine within an Availability Set.

[-update-domain <integer>] - Update Domain of This Instance

A cloud provider-assigned numerical update domain ID for this virtual machine within an Availability Set.

[-provider <text>] - Provider on which this instance is running.

The provider on which this instance is running.

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
          Version:  2010-08-31
Availability Zone:  us-east-1d
       Primary IP:  192.168.0.1
         Provider:  AWS

cluster1::> system node virtual-machine instance show
             Node:  node1
      Instance ID:  2831c724-97ca-4395-b8d3-a65c2a65b502._cloudontap-vm
    Deployment ID:  2831c724-97ca-4395-b8d3-a65c2a65b502
          Version:  2012-11-30
 Availability Set:  Fault Domain:  0
                    Update Doamin: 0
       Primary IP:  192.168.0.1
         Provider:  Azure
```

system script commands

system script delete

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 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

- system script show

system script show

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.
```

system script start

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 system script show `-this-sessionyes` command to display files that are recording the current CLI session. Use the 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

- system script show
- system script stop

system script stop

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 system script start command. Use the system script show `-this-sessionyes` 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

- system script start
- system script show

system script upload

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 system script show command to display saved CLI sessions. Use the 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 {(ftp|http)://(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

- system script show
- system script start

system service-processor commands

system service-processor reboot-sp

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::>
```

system service-processor show

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::>
```

system service-processor api-service modify

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 SP API service configuration. The SP API is a secure network API that enables Data ONTAP to communicate with the Service Processor over the network.

Parameters

`[-is-enabled {true|false}]` - Is SP API Service Enabled (privilege: advanced)

This parameter enables or disables the API service of the Service Processor. When the API service is disabled, features like network-based firmware updates and network-based down filer log collection will not be available, and the slower serial-interface will be used for firmware updates and down filer log collections.

`[-port <integer>]` - SP API Service Port (privilege: advanced)

This parameter specifies the port number on the Service Processor used for the API service. By default, port 50000 is used.

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 -port 50001

cluster1::*>system service-processor api-service show
 Service Processor API service configuration
is-enabled: true
      port: 50001
```

```
cluster1::*>system service-processor api-service modify -is-enabled false

cluster1::*>system service-processor api-service show
 Service Processor API service configuration
is-enabled: false
      port: 50001
```

system service-processor api-service renew-certificates

Renew SSL and SSH certificates used for secure communication with Service Processor API service

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

Description

The `system service-processor api-service renew-certificates` command renews the internal SSL and SSH certificates used for secure communication with the Service Processor API service. If the parameter `-renew-all` is not specified, only host certificates are renewed.

Parameters

`[-renew-all {true|false}]` - Renew CA Certificates Also (privilege: advanced)

This parameter specifies the type of certificates that needs to be renewed. If this parameter is set to false, only the host certificates (that is, the client and server certificates) are renewed. If this parameter is set to true, the root-ca certificate is renewed along with the host certificates.

Examples

The following example renews the host certificates and the root-ca certificates. The second command renews only the host certificates.

```
cluster1::*>system service-processor api-service renew-certificates -renew
-all true

cluster1::*>system service-processor api-service renew-certificates
```

system service-processor api-service show

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 API service configuration.

Examples

The following example displays the SP API service configuration:

```
cluster1::*>system service-processor api-service show
 Service Processor API service configuration
is-enabled: true
      port: 50000
```

system service-processor image modify

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.
```

system service-processor image show

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::>
```

system service-processor image update

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 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 Service Processor and how.

You can use the command system service-processor image update-progress show to check the progress of the update.

The following parameter combinations are not supported for this command:

- `-update-type` *differential* with `-clear-logs` *true*
- `-baseline` *true* with `-package` *<text>*

Parameters

-node {<nodename>|local} - Node

This parameter specifies the node whose Service Processor's 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 system node image package show command. If you do not specify this parameter, the Service Processor 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 Service Processor firmware version that is bundled with the currently running release of Data ONTAP. This is a safety mechanism that allows you to revert the SP firmware to the version that was qualified and bundled with the currently running version of Data ONTAP on your system. If not specified, this parameter defaults to false.

-update-type {serial-full|serial-differential|network-full} - Type

This parameter specifies the type of update to be performed.

If you set the value to *serial-full* , the command transfers contents of the entire SP firmware image to the SP via the packetized serial interface, and the contents are written to the SP primary partition.

If you set the value to *serial-differential* , the command transfers the SP firmware files that are different between the old SP firmware image and the new SP firmware image to the SP via the packetized serial interface, and the contents are written to SP primary partition.

If you set the value to *network-full* , the command transfers the entire SP firmware image to the SP via the SP network interface, and the contents are written to the SP primary partition.

If you do not specify the update-type option, the command checks for dependencies of network-based SP firmware updates. If the dependencies are satisfied, then the command uses the SP network interface to update the SP firmware. If the dependencies are not satisfied, then the serial interface is used.

[-clear-logs {true|false}] - Clear Logs After Update

If you set this parameter to *true* , the command resets log settings to factory default and clears contents of all logs maintained by the Service Processor, including:

- Event logs
- IPMI logs
- Forensics logs

Examples

The following command reverts the firmware on the Service Processor of the local node to the version that was packaged with the currently running release of Data ONTAP. A complete install will be performed, clearing all logs maintained by the Service Processor. The second command displays the status of the in-progress firmware install.

```
cluster1::> system service-processor image update -node local -update-type
full -baseline true -clear-logs 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

- system service-processor image update-progress show
- system node image package show

system service-processor image update-progress show

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 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 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 Service Processor 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 Service Processor 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 Service Processor 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 Service Processor 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 Service Processor 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 Service Processors of all nodes in the system.

```
cluster1::> system service-processor image update -node local -update-type
full -baseline true -clear-logs 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::>
```

system service-processor log show-allocations

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::>
```

system service-processor network modify

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}]` - 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::>
```

system service-processor network show

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}] - 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::>
```

system service-processor network auto-configuration disable

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
```

system service-processor network auto-configuration enable

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
```

system service-processor network auto-configuration show

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       -
```

system service-processor ssh add-allowed-addresses

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::>
```

system service-processor ssh remove-allowed-addresses

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::>
```

system service-processor ssh show

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::>
```

system services commands

system services firewall modify

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` .

`[-logging {true|false}]` - (DEPRECATED)-Enable Logging

This parameter is deprecated and may be removed in a future version of Data ONTAP.

Use this parameter to specify whether logging is enabled ("*`true`*") or disabled ("*`false`*") for the firewall service. The default setting is `false` .

Examples

The following example enables firewall protection and logging for a node named node1:

```
cluster1::> system services firewall modify -node node1 -enabled true
-logging true
```

system services firewall show

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 and logging information. 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*").

[-logging {true|false}] - (DEPRECATED)-Enable Logging

This parameter is deprecated and may be removed in a future version of Data ONTAP.

Selects information about the nodes with firewall logging 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 Logging
-------------- ------- -------
node0          true    false
node1          true    false
node2          true    false
node3          true    false
4 entries were displayed.
```

system services firewall policy clone

Clone an existing firewall policy

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

Description

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
```

system services firewall policy create

Create a firewall policy entry for a network service

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

Description

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, or 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 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.

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:

- default - The default protocol or protocols for the port to which the firewall is applied
- http - The HTTP protocol
- https - The HTTPS protocol
- ntp - The NTP protocol
- rsh - The RSH protocol
- snmp - The SNMP protocol
- ssh - The SSH 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 SSH 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
ssh -allow-list 192.0.2.128/25
```

The following example adds an entry to the firewall policy named data, associating the HTTPS 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
https -allow-list 192.0.2.128/25
```

Related Links

- network interface modify

system services firewall policy delete

Remove a service from a firewall policy

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

Description

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 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

- network interface modify

system services firewall policy modify

Modify a firewall policy entry for a network service

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

Description

The 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 SSH protocol to enable access from all addresses on the 192.0.2.128 subnet:

```
cluster1::> system services firewall policy modify -policy data -service
ssh -allow-list 192.0.2.128/25
```

Related Links

- system services firewall modify

system services firewall policy show

Show firewall policies

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

Description

The `system services firewall policy show` command displays information about 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
                     https      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
                     ssh        0.0.0.0/0, ::/0
```

system services manager install show

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.
```

system services manager policy add

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
```

system services manager policy remove

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
```

system services manager policy setstate

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 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

- system services manager policy show

system services manager policy show

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 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 (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

- system services manager status show

system services manager status show

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 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 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

- system services manager policy show

system services ndmp kill-all

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. This command is not supported on Infinite Volumes.

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
```

system services ndmp kill

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. This command is not supported on Infinite Volumes.

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
```

system services ndmp modify

(DEPRECATED)-Modify NDMP service configuration

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

Description

This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "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

This command is not supported on Infinite Volumes.

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

- vserver services ndmp modify

system services ndmp off

(DEPRECATED)-Disable NDMP service

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

Description

This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "vserver services ndmp off " command.

The `system services ndmp off` command is used to disable the NDMP service on any node in the cluster. This command is not supported on Infinite Volumes.

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

- vserver services ndmp off

system services ndmp on

(DEPRECATED)-Enable NDMP service

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

Description

> This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "vserver services ndmp on " command.

The `system services ndmp on` command is used to enable the NDMP service across any node in the cluster. This command is not supported on Infinite Volumes.

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

- vserver services ndmp on

system services ndmp password

(DEPRECATED)-Change the NDMP password for the node

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

Description

> This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "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. This command is not supported on Infinite Volumes.

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

- vserver services ndmp generate-password

system services ndmp probe

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

This command is not supported on Infinite Volumes.

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
  ....
   ...
```

system services ndmp show

(DEPRECATED)-Display NDMP service configuration

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

Description

This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "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. This command is not supported on Infinite Volumes.

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

- vserver services ndmp show

system services ndmp status

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. This command is not supported on Infinite Volumes.

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.
```

system services ndmp log start

(DEPRECATED)-Start logging for the specified NDMP session

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

Description

> This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "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 (privilege: advanced)

This parameter specifies the node.

`-session-id {<integer>|server}` - Session Identifier (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 (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

- vserver services ndmp log start

system services ndmp log stop

(DEPRECATED)-Stop logging for the specified NDMP session

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

Description

> This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data ONTAP. Use the Vserver-aware "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 (privilege: advanced)

This parameter specifies the node.

`-session-id {<integer>|server}` - Session Identifier (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

- vserver services ndmp log stop

system services ndmp node-scope-mode off

(DEPRECATED)-Disable NDMP node-scope-mode

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

Description

This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data 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.
```

system services ndmp node-scope-mode on

(DEPRECATED)-Enable NDMP node-scope-mode

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

Description

This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data 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 DATA 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.
```

system services ndmp node-scope-mode status

(DEPRECATED)-Status of NDMP node-scope-mode

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

Description

> This node-scoped NDMP command is deprecated. Node-scoped NDMP functionality may be removed in a future release of Data 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.
```

system services ndmp service modify

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

This command is not supported on Infinite Volumes.

Parameters

`-node {<nodename>|local}` - Node (privilege: advanced)

This specifies the node whose NDMP configuration is to be modified.

`[-common-sessions <integer>]` - NDMP Common Sessions (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.

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
```

system services ndmp service show

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. This command is not supported on Infinite 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.

[-node {<nodename>|local}] - Node (privilege: advanced)

Selects information about the specified node.

[-common-sessions <integer>] - NDMP Common Sessions (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.
```

system services ndmp service start

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 system services ndmp on command. The 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 system services ndmp service stop command. This command is not supported on Infinite Volumes.

Parameters

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

- system services ndmp on
- system services ndmp service stop

system services ndmp service stop

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 system services ndmp off command. The system services ndmp off command disables new NDMP connections on the node but does not stop the NDMP service daemon. This command is not supported on Infinite Volumes.

Parameters

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

- system services ndmp off

system services ndmp service terminate

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 system services ndmp kill-all for a clean termination of all active sessions on a node. This command is not supported on Infinite Volumes.

Parameters

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

- system services ndmp kill-all

system services web modify

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 (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 (privilege: advanced)

Defines whether HTTP is enabled. The default value for this parameter is *`false`*.

Examples

The following command changes the maximum size of the wait queue:

```
cluster1::> system services web modify -wait-queue-capacity 256
```

system services web show

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 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

- system services web node show

system services web node show

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 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 (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.

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

- system services web show

system smtape commands

system smtape abort

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 system smtape backup or 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

- system smtape backup
- system smtape restore

system smtape backup

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 copy name to perform an SMTape backup operation. The Snapshot copy name specified is used as the base Snapshot copy. 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 copy created during this backup operation in order to use this Snapshot copy 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 copy 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 copy 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 copy 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
```

system smtape break

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
```

system smtape continue

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
```

system smtape restore

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 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 system smtape break to get read/write permissions on the volume. Restore to an Infinite Volume is not supported. 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

- system smtape backup

- system smtape break

system smtape showheader

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 copy 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 Snapshot copies present in this backup.
- Snapshot ID - is the Snapshot ID of the backup Snapshot.
- Snapshot Time - time at which the backup Snapshot copy was created.
- Snapshot Name - name of the Snapshot copy 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
```

system smtape status clear

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 copy 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.
```

system smtape status show

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

Selectsinformation 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 copy name which matches the specified Snapshot copy 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
```

system snmp commands

system snmp authtrap

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:
              -   -
```

system snmp contact

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:
              -   -
```

system snmp enable-snmpv3

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::*>
```

system snmp init

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:
  -   -
```

system snmp location

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:
    -   -
```

system snmp prepare-to-downgrade

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
```

system snmp show

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:
    -   -
```

system snmp community add

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
```

system snmp community delete

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 wish 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.
```

system snmp community show

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
```

system snmp traphost add

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
```

system snmp traphost delete

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
```

system snmp traphost show

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
```

system status commands

system status show

Display System Status

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

Description

The `system status show` command displays information about the status of objects in Data ONTAP. You can limit output to specific types of information and specific status in Data ONTAP, 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 wish to view.

Parameters

{ [-fields <fieldname>,…]

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

| [-instance] }

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

[-internal-name <text>] - Internal Full Name

Selects status of objects that match this parameter value.

[-name <text>] - Name

Selects status of objects that match this parameter value.

[-vserver-id <text>] - Vserver ID

Selects status of objects that match this parameter value.

[-cluster-id <text>] - Cluster ID

Selects status of objects that match this parameter value.

[-is-cluster-scope {true|false}] - Cluster Scope

Selects status of objects that match this parameter value.

[-status <text>] - Status Value

Use this parameter to display the status.

[-update-time <MM/DD/YYYY HH:MM:SS>] - Update Time

Use this parameter to display the status last update time.

Examples

The following example displays information about system status in ONTAP:

```
cluster1::> system status show
  Cluster:tfarrellnscluster-1
    Node:tfarrell-vsim1
      Hypervisor system name:vsimesxrtp060.gdl.englab.netapp.com
        Version:5.5.0
        Cpu count:20
        Cpu version:CPU Pkg/ID/Node: 0/0/0 Intel(R) Xeon(R) CPU E5-2680 v2
@ 2.80GH
        Cpu count (virtual):40
        ONTAP guests:5
        Total guests:7
        Memory configured:255.9GB
        Memory used:52.08GB
        Physical Interface count:2
        Storage controller:LSI Logic / Symbios Logic 2004 iMR ROMB
        Storage configured:20GB
        Virtual Interface count:3
        Guest name:tfarrell_vsim_nsc1
          Memory configured:16GB
```

```
          Cpu count:4
        Virtual Interface:vmk0
          adminStatus:up
          Mtu size:1500
          operStatus:up
          speed:unlimited
        Virtual Interface:vmk1
          adminStatus:up
          Mtu size:9000
          operStatus:up
          speed:unlimited
        Virtual Interface:vmk2
          adminStatus:up
          Mtu size:9000
          operStatus:up
          speed:unlimited
        Physical NIC:vmnic0
          adminStatus:up
          Mtu size:9000
          operStatus:up
          speed:9.77GBps
        Physical NIC:vmnic1
          adminStatus:up
          Mtu size:9000
          operStatus:up
          speed:9.77GBps
      System up time:468:11:40
      Traps:
          1.3.6.1.4.1.6876.50.101.0:TRAPv1 SNMPv1 'public'
enterprise=1.3.6.1.4.1.6876.4.1 agent_addr=10.226.10.220 generic_trap=6
specific_trap=4 time-stamp=604476800
 [0]: 1.3.6.1.4.1.6876.50.101.0=INTEGER 4
 [1]: 1.3.6.1.4.1.6876.50.102.0=OCTET STRING 78: 2f 76 6d 66 73 2f 76 6f
6c 75 6d 65 73 2f 32 34 64 35 38 39 61 65 2d 36 34 33 62 31 35 38 65 2f 76
73 69 6d 73 79 73 2f 76 73 69 6d 73 2f 76 73 69 6d 73 79 73 5f 76 73 69 6d
5f 76 73 69 6d 31 2f 44 61 74 61 4f 4e 54 41 50 2e 76 6d 78
 [2]: 1.3.6.1.4.1.6876.2.1.1.2.2815=OCTET STRING 18: 76 73 69 6d 73 79 73
5f 76 73 69 6d 5f 76 73 69 6d 31
          1.3.6.1.4.1.6876.50.101.0:TRAPv1 SNMPv1 'public'
enterprise=1.3.6.1.4.1.6876.4.1 agent_addr=10.226.10.40 generic_trap=6
specific_trap=4 time-stamp=1683800
 [0]: 1.3.6.1.4.1.6876.50.101.0=INTEGER 4
 [1]: 1.3.6.1.4.1.6876.50.102.0=OCTET STRING 79: 2f 76 6d 66 73 2f 76 6f
6c 75 6d 65 73 2f 61 37 63 31 66 37 61 61 2d 39 65 39 63 63 34 61 30 2f 74
66 61 72 72 65 6c 6c 2f 76 73 69 6d 73 2f 74 66 61 72 72 65 6c 6c 5f 76 73
69 6d 5f 6e 73 63 31 2f 44 61 74 61 4f 4e 54 41 50 2e 76 6d 78
```

```
 [2]: 1.3.6.1.4.1.6876.2.1.1.2.2093=OCTET STRING 18: 74 66 61 72 72 65 6c
6c 5f 76 73 69 6d 5f 6e 73 63 31
47 entries were displayed.
```

system timeout commands

system timeout modify

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
```

system timeout show

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)
```

template commands

template copy

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
```

template delete

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
```

template download

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 {(ftp|http)://(hostname|IPv4 Address|'['IPv6 Address']')…}` - URI of the template (privilege: advanced)

This parameter specifies the URI from which the template will be downloaded.

`[-name <template name>]` - Name of the template (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
```

template provision

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%)
```

template rename

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
```

template show-permissions

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
```

template show

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
```

template upload

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 (privilege: advanced)

This parameter specifies the name of the template.

-uri {(ftp|http)://(hostname|IPv4 Address|'['IPv6 Address']')…} - URI to upload the template (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

template parameter modify

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 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

- template provision

template parameter show

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 paramter 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 maxium 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
```

volume commands

volume autosize

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 will automatically grow to when it is out of space or the minimum size that it will shrink 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

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 autosize configuration.

`{ [-maximum-size {<integer>[KB|MB|GB|TB|PB]}]` - Maximum Autosize

This parameter allows the user to specify the maximum size to which a flexible volume can grow. The default for FlexVol volumes is 120% of the volume size. If the value of this parameter is invalidated by manually resizing the volume or is invalid when the autosize feature is enabled, 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 FlexVol volume size. If you specify a larger value, the value of `-maximum-size` is automatically reset to the supported maximum without returning an error. This parameter is not supported on Infinite Volumes.

`[-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. 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 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, `-mode` is *off* for new FlexVol 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. This parameter is not supported on Infinite Volumes.

| [-reset <true>] - Autosize Reset }

This option 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 based on the current size of the volume. For example, the max-autosize value will be 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.
```

volume create

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. Only Infinite Volumes can use unified.)
- 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 FlexVols
- Maximum size for autosizing FlexVols
- Minimum size for autosize
- Grow used space threshold percentage for autosize
- Shrink used space threshold percentage for autosize
- Whether autosizing is enabled for FlexVols
- 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
- Whether the volume create operation runs as a foreground or background process
- Caching policy
- Encrypt
- Cache retention priority
- Efficiency policy

The ability to create Infinite Volumes is deprecated and may be removed in a future release of Data ONTAP. If you are using Infinite Volumes it is recommended that you do not upgrade the cluster to a release that is later than Data ONTAP 9.3.0.

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 150 or fewer characters in length for Infinite Volumes, 197 or fewer characters in length for FlexGroups, 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. An aggregate may be specified multiple times to have multiple constituents created on it. This parameter only applies to FlexGroups.

[-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 FlexGroups

| -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.

This parameter only applies to FlexGroups.

[-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 FlexGroup volume is 200 MB per constituent. However, the recommended size for a FlexGroup volume is a minimum of 100 GB per constituent. For Infinite Volumes, the minimum size is 1.33 TB per node that will host a data 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 volume modify command. The maximum number of files a volume is configured for is listed under "Total Files" when running the command 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.

The *`mixed`* state applies to FlexGroups and Infinite 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 vserver export-policy create command. FlexGroups 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. The *`unified`* security style can only be used on Infinite Volumes.

`[-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 (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 (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 FlexGroups or Infinite 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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Directory Size (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. This parameter is not supported on Infinite Volumes. 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 Snapshot copies 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 Snapshot copies 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 or LUNs 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 Snapshot copies.

| [-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. Because CIFS does not handle out-of-space conditions, do not use the value *`none`* if the volume is accessible to CIFS clients. 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.

[-percent-snapshot-space <percent>] - Space Reserved for Snapshot Copies

This optionally specifies the amount of space that is reserved in the volume for Snapshot copies. 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 vserver create and vserver modify commands. The schedules associated with the `snapshot-policy` for an Infinite Volume cannot have an interval shorter than hourly. 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. You cannot specify the language encoding setting for an Infinite Volume.

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 FlexGroups and Infinite 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.

[-storage-service <storage service name>] - Storage Service Name (privilege: advanced)

The name of the initial storage service for the Infinite Volume. This is required if the parameter `-is-managed-by-service` is set to *true* . This parameter applies to Infinite Volumes only.

[-enable-snapdiff {true|false}] - Create Namespace Mirror Constituents For SnapDiff Use

When set to *true* for an Infinite Volume that spans three or more nodes, namespace mirror constituents are created for SnapDiff use. One namespace mirror constituent is created on every node that contains a data constituent for the Infinite Volume. A namespace constituent is not created on nodes that contain either the namespace constituent or a namespace mirror constituent used for data protection of the namespace constituent. An automatic daily replication schedule is set up for every namespace mirror constituent created. The default setting is *false* . This parameter applies to Infinite Volumes only.

[-unreachable-attr-action {return-generated|wait}] - Action When Attributes Are Not Reachable (privilege: advanced)

This parameter specifies the information that an Infinite Volume returns when a client lists a directory that contains one or more files with inaccessible attributes, which can happen when a data constituent is not

online. When this parameter is set to *return-generated*, the Infinite Volume returns default values for the attributes, which appear to the client as a file size of 0 and timestamps that are in the past. When this parameter is set to *wait*, the Infinite Volume returns a RETRY error, which may cause some clients to hang. When the inaccessible file attributes become available, the Infinite Volume returns them to the client. The default setting is *return-generated*. This parameter applies to Infinite Volumes only.

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

The name of the aggregate in which to create the Infinite Volume namespace constituent. If a name is not provided, Data ONTAP picks the aggregate assigned to the Vserver that has the most usable space. This parameter applies to Infinite Volumes only.

[-max-namespace-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Namespace Constituent (privilege: advanced)

The maximum size of the namespace constituent. The default value is 10TB. This parameter applies to Infinite Volumes only.

[-ns-mirror-aggr-list <aggregate name>,…] - List of Aggregates for Namespace Mirrors (privilege: advanced)

Specifies the aggregates that can be used to create Infinite Volume namespace mirror constituents. No other aggregate will be chosen for this purpose. Aggregates in this list will remain available for other uses in the Infinite Volume. This parameter applies to Infinite Volumes only.

[-data-aggr-list <aggregate name>,…] - List of Aggregates for Data Constituents (privilege: advanced)

Specifies the aggregates that can be used to create Infinite Volume data constituents. No other aggregate will be chosen for this purpose. Aggregates in this list will remain available for other uses in the Infinite Volume. This parameter applies to Infinite Volumes only.

[-max-data-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Each Data Constituent (privilege: advanced)

This optionally specifies the maximum size of an Infinite Volume data constituent. The default value is determined by checking the maximum FlexVol size setting on all nodes used by the Infinite Volume. The smallest value found is selected as the default for the `max-data-constituent-size` for the Infinite Volume. This parameter applies to Infinite Volumes only.

{ [-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. This parameter is not supported on Infinite Volumes.

| [-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. This parameter is not supported on FlexGroups or Infinite 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 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. This parameter is not supported on Infinite Volumes.

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 (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. This parameter is not supported on Infinite Volumes.

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 (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 (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 FlexGroups only.

[-efficiency-policy <efficiency policy>] - Storage Efficiency Policy (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 Infinite Volumes or data protection volumes on any platform. To disable compression on the volume in All-Flash FAS, use the value *none* . The default value is *inline-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. The default value for a volume of type “RW” is *protected* . This parameter is not supported on Infinite Volumes.

[-encrypt <true>] - Enable 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 will be encrypted using the generated key. By default, volume created is not encrypted.

[-tiering-policy {snapshot-only|none|backup}] - 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 capacity tier when they become cold. FabricPool combines flash (performance tier) with an object store (external capacity tier) into a single aggregate. If a tiering policy is not assigned to this volume, then 'snapshot-only' is assigned as the default policy. 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 Snapshot copies not associated with the active file system.
- none - Volume blocks will not be tiered to the capacity tier.
- backup - On DP volumes this policy allows all transferred user data blocks to start in the capacity tier.

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 Data ONTAP.

```
cluster1::> volume create -vserver vs0 -volume fg -auto-provision-as
flexgroup
```

Related Links

- volume modify
- volume show
- vserver export-policy create
- vserver create
- vserver modify

volume delete

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 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 (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

- volume offline

volume expand

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 FlexGroups.

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 Snapshot copies to be set to "partial". Partial Snapshot copies
      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 Snapshot copies to be set to "partial". Partial Snapshot copies
      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
```

volume make-vsroot

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 atleast 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`* .

If you run this command on a Vserver with Infinite Volume, you must create a new root volume. You cannot promote an existing non-root volume to be the root volume of the Vserver with Infinite Volume. When the command executes, the new root volume is created on the aggregate specified by the *`aggregate`* parameter, and the old root volume is deleted.

This command is available only at the advanced privilege level and higher.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

This specifies the Vserver on which a non-root volume is to be made the root volume.

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

This specifies the non-root volume that is to be made the root volume of its Vserver. For Vservers with FlexVol volume this must be an existing FlexVol volume. For Vservers with Infinite Volume this must be a non-existent volume that will be created during the execution of the command. Using a SnapLock volume as the root volume for a Vserver is not supported.

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

This specifies the aggregate that the new root volume will be created on for Vservers with Infinite Volume. This parameter is not supported for Vservers with FlexVol volumes.

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
```

volume modify

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. Only Infinite Volumes can use unified.)
- 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
- Convert ucode
- Caching policy
- Cache retention priority

You can use the `volume move` command to change a volume's aggregate or node. You can use the volume rename command to change a volume's name. You can use the 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. This parameter is not supported on Infinite Volumes that are managed by storage services.

[-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.

The *mixed* state applies to Infinite 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 vserver export-policy create command. FlexGroups 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. The *unified* security style can only be used on Infinite Volumes.

[-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 (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. This parameter is not supported on Infinite Volumes. 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. This parameter is not supported on Infinite Volumes. 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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

| [-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. This parameter is not supported on Infinite Volumes.

[-files <integer>] - Total Files (for user-visible data)

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.

[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Directory Size (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. This parameter is not supported on Infinite Volumes.

{ [-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. This parameter is not supported on Infinite Volumes. 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 Snapshot copies 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 Snapshot copies 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 or LUNs 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 Snapshot copies.

| [-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. This parameter is not supported on FlexGroups or Infinite Volumes that are managed by storage services.

[-fractional-reserve <percent>] - Fractional Reserve }

This option changes the amount of space reserved for overwrites of reserved objects (LUNs, files) in a volume. This parameter is not supported on Infinite Volumes. 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 (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 (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 Snapshot copies. 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 Vserver, as specified by the `-snapshot-policy` parameter of the vserver create and vserver modify commands. The schedules associated with the `snapshot-policy` for an Infinite Volume cannot have an interval shorter than hourly. When replacing a `snapshot-policy` on a volume, any existing Snapshot copies 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 Snapshot copies if the prefixes do not match. Once the new `snapshot-policy` takes effect, depending on the new retention count, any existing Snapshot copies that continue to use the same prefixes may be deleted. For example, if your existing `snapshot-poilcy` is setup to retain 150 weekly snapshot copies and you create a new `snapshot-policy` that uses the same prefixes but changes the retention count to 50 snapshot copies. Once the new `snapshot-policy` takes effect it will start deleting older snapshot copies until there are only 50 remaining.

[-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 Infinite 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 (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, and is not supported on FlexGroups or Infinite Volumes.

[-filesys-size-fixed {true|false}] - Is File System Size Fixed

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 traditional volume and the size is larger than the file system size, setting this option to false forces the file system to grow to the size of the 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

This parameter has been deprecated and may be removed in a future release of Data 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 Snapshot copies into the active file system, and
will result in conservative space utilization. Using `space-optimized`
may be useful when the volume has Snapshot copies 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 Snapshot
copies. 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 Snapshot copies 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 Snapshot copies to reclaim space. If the option is set to snap_delete the system will first automatically delete Snapshot copies and in case of failure to reclaim space will try to grow the volume. This parameter is not supported on Infinite Volumes.

[-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 Snapshot copies into the active file system, and will result in conservative space utilization. Using `space-optimized` may be useful when the volume has Snapshot copies 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 Snapshot copies. 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 Snapshot copies. If set to *create-time* , automatic Snapshot copies 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 copy is named using the format *<schedule_name>.<n>* . Example: "hourly.0". Older automatic Snapshot copies are named using the format *<schedule_name>.yyyy-mm-dd_hhmm* . Example: "hourly.2010-04-01_0831".

[-storage-service <storage service name>] - Storage Service Name (privilege: advanced)

The name of the storage service for the Infinite Volume. This parameter applies to Infinite Volumes only.

[-enable-snapdiff {true|false}] - Create Namespace Mirror Constituents For SnapDiff Use

When set to *true* for an Infinite Volume that spans three or more nodes, namespace mirror constituents are created for SnapDiff use. One namespace mirror constituent is created on every node that contains a data constituent for the Infinite Volume. A namespace mirror constituent is not created on nodes that contain either the namespace constituent or a namespace mirror constituent used for data protection of the namespace constituent. An automatic daily replication schedule is set up for every namespace mirror constituent created. If set to *false* , all existing namespace mirror constituents used by SnapDiff are deleted. The namespace mirror constituent used for namespace data protection is not affected. This parameter applies to Infinite Volumes only.

[-unreachable-attr-action {return-generated|wait}] - Action When Attributes Are Not Reachable (privilege: advanced)

This parameter specifies the information that an Infinite Volume returns when a client lists a directory that contains one or more files with inaccessible attributes, which can happen when a data constituent is not online. When this parameter is set to *return-generated* , the Infinite Volume returns default values for the attributes, which appear to the client as a file size of 0 and timestamps that are in the past. When this parameter is set to *wait* , the Infinite Volume returns a RETRY error, which may cause some clients to hang. When the inaccessible file attributes become available, the Infinite Volume returns them to the client. The default setting is *return-generated* . This parameter applies to Infinite Volumes only.

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

The name of the aggregate in which to create the Infinite Volume namespace constituent. If a name is not provided, Data ONTAP picks the aggregate assigned to the Vserver that has the most usable space. This parameter can only be specified for DP Infinite Volumes that do not have an existing namespace constituent. This parameter applies to Infinite Volumes only.

[-max-namespace-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Namespace Constituent (privilege: advanced)

The maximum size of the namespace constituent. The default value is 10TB. This parameter applies to Infinite Volumes only.

[-ns-mirror-aggr-list <aggregate name>,…] - List of Aggregates for Namespace Mirrors (privilege: advanced)

Specifies the aggregates that can be used to create Infinite Volume namespace mirror constituents. No

other aggregate will be chosen for this purpose. Aggregates in this list will remain available for other uses in the Infinite Volume. This parameter applies to Infinite Volumes only.

[-max-data-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Each Data Constituent (privilege: advanced)

This parameter specifies the maximum size of an Infinite Volume data constituent. The default value is determined by checking the maximum FlexVol size setting on all nodes used by the Infinite Volume. The smallest value found is selected as the default for the maximum constituent size. This parameter applies to Infinite Volumes only.

{ [-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". This parameter is not supported on Infinite Volumes.

| [-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". This parameter is not supported on FlexGroups or Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

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 (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 (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 FlexGroups 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. This parameter is not supported on Infinite Volumes.

[-tiering-policy {snapshot-only|none|backup}] - 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 capacity tier when they become cold. FabricPool combines flash (performance tier) with an object store (external capacity tier) into a single aggregate. 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 Snapshot copies not associated with the active file system.
- none - Volume blocks will not be tiered to the capacity tier.
- backup - On DP volumes this policy allows all transferred user data blocks to start in the capacity tier.

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

- volume rename
- volume make-vsroot
- vserver export-policy create
- vserver create
- vserver modify

volume mount

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
```

volume offline

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 may not be taken offline. A number of operations being performed on the volume in question can prevent volume offline from succeeding for various lengths of time. If such operations are required, the command may take additional time to complete. If they do not, 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 Infinite Volumes. For FlexVol volumes, the command always runs in the foreground.

[-disable-luns-check <true>] - Disable Check for Existing LUNs

Taking the volume offline will make all associated LUNs 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.
```

volume online

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 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 Infinite 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.
```

volume rehost

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.

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.

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
```

volume rename

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 on the same Vserver. The `volume rename` command is not supported for FlexGroups.

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 Infinite 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
```

volume restrict

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 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 Infinite 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

- volume offline

volume show-footprint

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. + + The `volume show-footprint` command is not supported for Infinite Volumes. However, the command displays information about Infinite Volume constituents as if the constituents were FlexVol 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 Snapshot copies. This row only includes data and not reserved space, so when volumes have reserved files, the volume's total used in the volume show-space command output can exceed the value in this row.

[-volume-blocks-footprint-percent <percent>] - 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 Data 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>] - 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 an external capacity 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 external capacity 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 capacity 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 external capacity 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 capacity tier.

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 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 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 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 an external capacity 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 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 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 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 Footprint                           1.25GB      13%

4 entries were displayed.
```

Related Links

- volume show-space

volume show-space

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. + + The `volume show-space` command is not supported for Infinite Volumes; however, the command displays information about Infinite Volume constituents as if the constituents were FlexVol 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 metadata (for example indirect blocks, directory blocks) directly associated with user files 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 accomodate 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.

This is equivalent to the used field in the output of the 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 differs from *total-used* space by the space that is reserved for future writes. The value includes blocks in use by Snapshot copies.

[-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 Snapshot copies.

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 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 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 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 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 Used                                193.3MB        79%
Total Physical Used                       192.9MB        79%
2 entries were displayed.
```

Related Links

- volume show

volume show

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, flex or infinitevol)
- 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
- 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 Snapshot Copies
- Percentage of Snapshot copy 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 Snapshot Copies
- Caching policy
- FlexGroup master data set ID
- FlexGroup index
- FlexGroup UUID
- Maximum size of the FlexGroup constituent
- Whether the volume has FlexGroup enabled
- Whether a FlexGroup is Qtree enabled
- List of the aggregates used by the FlexGroup
- List of the nodes used by the FlexGroup
- SnapLock Type
- Is in pre-commit phase of Copy-Free Transition
- Application that the volume belongs to
- Application UUID

To display detailed information about all volumes, run the command with the `-instance` parameter. Fields not supported by Infinite Volumes will display a value of "**-**".

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] (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 Snapshot copies 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 FlexGroups.

[-aggr-list <aggregate name>,…] - List of Aggregates for FlexGroup Constituents

If this parameter is specified, the command displays information only about the FlexGroup or FlexGroups that are located on the specified list of storage aggregates. This parameter applies to FlexGroups 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, Snapshot copies, 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 (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 FlexGroups.

[-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 FlexGroups and Infinite Volumes. If the state of a FlexGroup or Infinite 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|infinitevol}] - 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, and *infinitevol* for Infinite Volumes.

[-volume-style-extended {flexvol|infinitevol|infinitevol-constituent|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, *infinitevol* for Infinite Volumes, *infinitevol-constituent* for Infinite Volume constituents, *flexgroup* for FlexGroups and *flexgroup-constituent* for FlexGroup constituents.

[-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 or Infinite Volume, depending on the value provided.

[-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 (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.

It includes Snapshot spill (the amount of space by which Snapshot copies 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 Snapshot copies 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 FlexGroups.

[-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 FlexGroups.

[-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 FlexGroups.

[-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 FlexGroups.

[-files <integer>] - Total Files (for user-visible data)

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>] - Files Used (for user-visible data)

If this parameter is specified, the command displays information only about the volume or volumes that have the specifies number of files used.

[-maxdir-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Directory Size (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 (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 (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-copy 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 Snapshot copies.

[-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 Snapshot copies.

[-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 FlexGroups.

[-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 (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).

[-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

This parameter has been deprecated and may be removed in a future release of Data 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 FlexGroups.

[-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-copy 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 (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 (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 (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 (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. Infinite Volumes will report the aggregated setting of their constituent data volumes as *true* or *false* if all constituents have the same setting, otherwise no value will be reported.

[-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

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]}] - Space 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 Snapshot copies.

[-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}] - 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.

[-is-managed-by-service {true|false}] - Managed By Storage Service

If this parameter is specified as `true` , only display volumes that are managed by storage services. This field is displayed as "-" for FlexGroups.

[-storage-service <storage service name>] - Storage Service Name (privilege: advanced)

If this parameter is specified, the command displays information only about volumes that match the specified storage service.

[-enable-snapdiff {true|false}] - Create Namespace Mirror Constituents For SnapDiff Use

Setting this parameter displays information only about Infinite Volumes that either do or do not have namespace mirror constituents for SnapDiff use, depending on the value provided. This parameter applies to Infinite Volumes only.

[-unreachable-attr-action {return-generated|wait}] - Action When Attributes Are Not Reachable (privilege: advanced)

This parameter specifies the information that an Infinite Volume returns when a client lists a directory that contains one or more files with inaccessible attributes. If this parameter is specified, the command displays information only about volumes that match the specified action. This parameter is not supported for FlexVol volumes.

[-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 Infinite Volumes and FlexGroups only.

[-max-namespace-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Namespace Constituent (privilege: advanced)

If this parameter is specified, the command displays information only about volumes that match the specified namespace constituent size.

[-max-data-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum Size of Each Data Constituent (privilege: advanced)

If this parameter is specified, the command displays information only about the Infinite Volume or Infinite Volumes that have the specified maximum data constituent size. This parameter applies to Infinite Volumes only.

[-is-cft-precommit {true|false}] - Is in the precommit phase of Copy-Free Transition (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 (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. This parameter is not supported on Infinite Volumes. 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 FlexGroups.

[-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 Snapshot copies.

[-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 (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 FlexGroups.

[-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. This differs from *total-used* space by the space that is reserved for future writes. The value includes blocks in use by Snapshot copies.

[-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 Snapshot copies.

[-flexgroup-msid <integer>] - FlexGroup Master Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup or FlexGroup constituents that have the specified FlexGroup master data-set ID. This parameter applies to FlexGroups and FlexGroup constituents only.

[-flexgroup-index <integer>] - FlexGroup Index (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup constituents that have the specified FlexGroup index. This parameter applies to FlexGroup constituents only.

[-flexgroup-uuid <UUID>] - UUID of the FlexGroup (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup or FlexGroup constituents that have the specified FlexGroup UUID. This parameter applies to FlexGroups and FlexGroup constituents only.

[-max-constituent-size {<integer>[KB|MB|GB|TB|PB]}] - Maximum size of a FlexGroup Constituent (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroup or FlexGroups that have the specified maximum constituent size. This parameter applies to FlexGroups only.

[-inofile-version <integer>] - Inofile Version (privilege: advanced)

If this parameter is specified, the command displays information only about the volumes whose inode files are at the specified version.

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

If this parameter is specified, the command displays information only about the FlexGroup or FlexGroups that are located on the specified list of storage systems. This parameter applies to FlexGroups only.

[-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 FlexGroups or not, depending on the value provided.

[-is-flexgroup-qtree-enabled {true|false}] - Is FlexGroup Qtree Support Enabled (privilege: advanced)

If this parameter is specified, the command displays information only about the FlexGroups with Qtrees enabled or disabled, depending on the value provided.

[-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. This parameter is not supported on Infinite Volumes.

[-encrypt <true>] - Enable Encryption

If this parameter is specified, the command displays information only about the volumes that are encrypted.

[-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 not encrypted (false).

[-encryption-state {none|full|partial|converting_to_encrypted|converting_to_plaintext|rekeying}] - Volume 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* , and *partial* . The value *partial* is used for FlexGroups, 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.

[-application <text>] - Application

Selects the volumes that are part of an application that matches the parameter value.

[-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 not available in the aggregate. This applies only for a 'none' guaranteed volume when it's unused size exceeds available space in the aggregate. This value is always zero for 'volume' guaranteed volumes. This parameter is not supported on Infinite Volumes.

[-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 Snapshot copies 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. This parameter is not supported on Infinite Volumes.

[-tiering-policy {snapshot-only|none|backup}] - 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 capacity tier when they become cold. FabricPool combines flash (performance tier) with an object store (external capacity tier) into a single aggregate. 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 Snapshot copies not associated with the active file system.
- none - Volume blocks will not be tiered to the capacity tier.
- backup - On DP volumes this policy allows all transferred user data blocks to start in the capacity tier.

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 flexvols only): 8.40GB
                                Minimum Autosize: 30MB
             Autosize Grow Threshold Percentage: 85%
               Autosize Shrink Threshold Percentage: 50%
                                Autosize Mode: off
               Autosize Enabled (for flexvols 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: -
Create Namespace Mirror Constituents For SnapDiff Use: -
                            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 Master Data Set ID: -
                                  FlexGroup Index: -
                             UUID of the FlexGroup: -
          Maximum size of a FlexGroup 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
```

volume size

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 <text>]` - [+|-]<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. This parameter is not supported on Infinite Volumes that are managed by storage services.

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.
```

volume transition-prepare-to-downgrade

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 clustered Data 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 clustered Data 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.

volume unmount

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 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

- volume mount

volume aggregate commands

volume aggregate vacate

Move all Infinite Volume constituents from one aggregate to another.

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

Description

The `volume aggregate vacate` command moves all constituents belonging to a given Infinite Volume from the source aggregate to the destination aggregate.

The `volume aggregate vacate` command is only supported for Infinite Volumes.

Parameters

`-vserver <vserver name>` - Vserver Name

The name of the Vserver that owns the volume.

`-volume <volume name>` - Volume Name

The name of the volume.

`-source-aggregate <aggregate name>` - Source Aggregate

The source aggregate from which all Infinite Volume constituents are being moved.

`-destination-aggregate <aggregate name>` - Destination Aggregate

The destination aggregate to which the Infinite Volume constituents are being moved.

`[-w, -foreground {true|false}]` - Foreground Process

This specifies whether the operation runs as a foreground process. If this parameter is not specified, the default setting is *`false`* (the operation runs in the background). When set to *`true`* , the command will not return until the process completes.

Examples

The following example moves all constituents of Infinite Volume repo_vol from aggr1 to aggr2.

```
cluster1::> volume aggregate vacate -vserver vs0 -volume repo_vol -source
-aggregate aggr1 -destination-aggregate aggr2
```

volume clone commands

volume clone create

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 only supported for flexible 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

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 FlexGroups or Infinite 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. This parameter is not supported on Infinite Volumes. 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. This parameter is not supported on Infinite Volumes.

[-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.

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
```

volume clone show

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.

[-inodes-processed <integer>] - Inodes Processed

Selects summary information for the FlexClone volumes that have the specified number of Inodes processed for splitting the FlexClone volume from its parent volume.

[-inodes-total <integer>] - Total Inodes

Selects summary information for the FlexClone volumes that have the specified number of total Inodes.

[-inode-percentage-complete <integer>] - Inode Percentage Complete

Selects summary information for the FlexClone volumes that have specified percentage of Inodes processed for splitting the FlexClone volume from its parent volume.

[-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. This parameter is not supported on Infinite Volumes. 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 (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>] - Block 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.

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.
```

volume clone split estimate

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.

[-inodes-processed <integer>] - Inodes Processed

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified number of Inodes have been processed already.

[-inodes-total <integer>] - Total Inodes

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified total number of inodes need to be processed.

[-inode-percentage-complete <integer>] - Inode Percentage Complete

This specifies the free disk space estimates for splitting the FlexClone volumes for which the specified percentage of Inode processing has been completed.

[-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 greeements (SLAs) that adjust based on the volume allocated space or used space. This parameter is not supported on FlexGroups or Infinite 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 (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. This parameter is not supported on Infinite Volumes.

[-block-percentage-complete <integer>] - Block 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.

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.
```

volume clone split show

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
- Number of inodes processed during clone splitting
- Total inodes to be processed during clone splitting
- Percentage of inodes processed
- 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.

[-inodes-processed <integer>] - Inodes processed

Selects information about all the ongoing FlexClone splitting jobs which have the specified number of Inodes processed.

[-inodes-total <integer>] - Total Inodes

Selects information about all the ongoing FlexClone splitting jobs that have the specified number of total Inodes to be processed.

[-inode-percentage-complete <integer>] - Inode Percentage complete

Selects information about all the ongoing FlexClone splitting jobs that have the specified percentage of Inode processing completed.

[-block-percentage-complete <integer>] - Block 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)
                                Inodes              Blocks
                          --------------------- ---------------------
Vserver   FlexClone      Processed      Total Scanned Updated % Inode  %
Block
                                    Complete Complete
--------- ------------- ---------- ---------- ------- ------- --------
---------
vs1       fc_vol_1               0       1260     0     0         0   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
         FlexClone Volume: fc_vol_2
         Inodes processed: 0
             Total Inodes: 3192
Inode Percentage complete: 0
Block Percentage complete: 0
           Blocks Scanned: 0
           Blocks Updated: 0
```

volume clone split start

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 volume clone split stop command. You can also stop the job using the job stop command. You can monitor the current progress of the job using the volume clone split show and 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`.

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

- volume clone split stop
- job stop
- volume clone split show
- job show

volume clone split stop

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
```

volume efficiency commands

volume efficiency check

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 FlexGroups or Infinite Volumes that are managed by storage services.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)
Specifies the Vserver on which the volume is located.

{ `-volume <volume name>` - Volume Name (privilege: advanced)
Specifies the volume on which the verify operation needs to be started.

| `-path </vol/volume>` - Volume Path (privilege: advanced) }
Specifies the volume path on which the verify operation needs to be started.

`[-d, -delete-checkpoint {true|false}]` - Delete Checkpoint (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
```

volume efficiency modify

Modify the efficiency configuration of a volume

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

Description

This command is used to set or modify the schedule, policy and various other efficiency configuration options on a volume. This command is not supported on Infinite Volumes that are managed by storage services.

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`* policy when there is an active background operation on the volume.

Note that schedule and policy are mutually exclusive options.

`[-compression-type {none|secondary|adaptive}]` - Compression Type (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.

[-compression {true|false}] - Compression

This option is used to enable and disable compression. The default value is false.

[-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 false.

You can use the *`inline-only`* predefined efficiency policy to run inline compression without the need of any background efficiency operations.

[-inline-dedupe {true|false}] - Inline Dedupe

This option is used to enable and disable inline deduplication. The default value is false.

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 false.

[-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 false.

[-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 false.

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
```

volume efficiency off

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. This command is not supported on Infinite Volumes that are managed by storage services.

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
```

volume efficiency on

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 volume efficiency modify command can be used to modify schedule and the volume efficiency policy modify command can be used to modify policy. You can also manually start an efficiency operation with the volume efficiency start command. This command is not supported on Infinite Volumes that are managed by storage services.

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
```

```
cluster1::> volume efficiency on -vserver vs1 -volume vol1 -needs-upgrade
true
```

```
cluster1::> volume efficiency on -vserver vs1 -path /vol/vol1 -needs
-upgrade true
```

Related Links

- volume efficiency modify
- volume efficiency policy modify
- volume efficiency start

volume efficiency prepare-to-downgrade

Identify any incompatable 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 releases prior to ONTAP 9. This command also disables the use of incompatible efficiency features. This command is not supported on FlexGroups.

Parameters

[-disable-feature-set <downgrade version>] - Data ONTAP Version (privilege: advanced)

This parameter specifies the Data ONTAP version that introduced new volume efficiency feature set.

Examples

The following example disables the the features introduced in Data ONTAP 8.3.1

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.1
```

The following example disables the the features introduced in Data ONTAP 8.3.2.

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.2
```

The following example ignores offline volumes while disabling the the features introduced in Data ONTAP 8.3.2
.

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.2 -skip-offline-volumes true
```

The following example ignores offline volumes while disabling the the features introduced in Data ONTAP 8.3.1
.

```
cluster1::*> volume efficiency prepare-to-downgrade -disable-feature-set
8.3.1 -skip-offline-volumes true
```

volume efficiency promote

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. This command is not supported on Infinite Volumes that are managed by storage services.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ `-volume <volume name>` - Volume Name (privilege: advanced)

This specifies the name of the volume on which auto scheduling needs to be restarted.

| `-path </vol/volume>` - Volume Path (privilege: advanced) }

This specifies the volume path on which auto scheduling needs to be restarted.

Examples

The following examples promote a volume from deprioritized state back to auto state.

```
cluster1::> volume efficiency promote -vserver vs1 -volume vol1
```

volume efficiency revert-to

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 Data ONTAP. This command is not supported on FlexGroups.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ `-volume <volume name>` - Volume Name (privilege: advanced)

This specifies the name of the volume for which volume efficiency metadata needs to be reverted.

| `-path </vol/volume>` - Volume Path (privilege: advanced) }

This specifies the volume path for which volume efficiency metadata needs to be reverted.

`[-version <revert version>]` - Revert to Version (privilege: advanced)

Specifies the version of Data ONTAP to which the volume efficiency metadata needs to be formatted.

`[-d, -delete {true|false}]` - Delete Existing Metafile on Revert (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 (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 (privilege: advanced)

If set to `true` , this parameter Specifies that the volume efficiency metadata needs to be reverted to minor version of Data ONTAP. By default this parameter is set to `false` .

`[-check-snapshot {true|false}]` - Revert ignore snapshots (privilege: advanced)

If set to *`false`* , this parameter specifies that the volume efficiency revert will not check for Snapshot copies created by previous releases of Data 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 8.3.

```
cluster1::> volume efficiency revert-to -vserver vs1 -volume vol1 -version
8.3
```

```
cluster1::> volume efficiency revert-to -vserver vs1 -path /vol/vol1
-version 8.3
```

volume efficiency show

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 Data ONTAP release is active.
- *`Disabled`* : Efficiency is disabled on the volume.

```
 Status is not supported for Infinite Volumes and will display a value
of "-"
* 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". Progress is not supported for Infinite Volumes
and will display a value of "-"
```

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 (fields not supported by Infinite Volumes will display a value of "-"):

- 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. This parameter is not supported on Infinite Volumes.
- Last Operation State: Status of the last operation (Success or Failure). Not supported on Infinite Volumes.
- Last Successful Operation Begin: The time and date at which the last successful operation began. This parameter is not supported on Infinite Volumes.
- Last Successful Operation End: The time and date at which the last successful operation ended. This parameter is not supported on Infinite Volumes.
- Last Operation Begin: The time and date at which the last operation began. This parameter is not supported on Infinite Volumes.
- Last Operation End: The time and date at which the last operation ended. This parameter is not supported on Infinite Volumes.
- Last Operation Size: The size of the last operation. This parameter is not supported on Infinite Volumes.
- Last Operation Error: The error encountered by the last operation. This parameter is not supported on Infinite Volumes.
- Change Log Usage: The percentage of the change log that is used. This parameter is not supported on Infinite Volumes.
- Logical Data: The total logical data in the volume, and how much is reached compared to the deduplication logical data limit. This parameter is not supported on Infinite Volumes.
- 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 Data ONTAP release is queued.
- Stale Fingerprints: The percentage of stale entries in the fingerprint database. If this is greater than 20 percent a subsequent volume efficiency start operation triggers the verify operation, which might take a long time to complete. This parameter is not supported on Infinite Volumes.
- 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).

- Inline Adaptive Data Compaction: Whether Inline Adaptive Data Compaction is enabled or disabled on the volume. When enabled, Data 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.

No information is displayed for Infinite Volumes that are managed by storage services.

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. This parameter is not supported on Infinite Volumes.

[-progress <text>] - Progress

Displays information only for those volumes that match the specified progress. This parameter is not supported on Infinite Volumes.

[-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.

[-compression-type {none|secondary|adaptive}] - Compression Type (privilege: advanced)

Displays information about the type of compression on the volume[adaptive or secondary].

[-blks-skipped-sharing <integer>] - Blocks Skipped Sharing

Displays information only for those volumes that match the specified blocks skipped sharing. This parameter is not supported on Infinite Volumes.

[-last-op-state <text>] - Last Operation State

Displays information only for those volumes that match the specified last operation state. This parameter is not supported on Infinite Volumes.

[-last-success-op-begin <Date>] - Last Success Operation Begin

Displays information only for those volumes that match the specified last successful operation begin time. This parameter is not supported on Infinite Volumes.

[-last-success-op-end <Date>] - Last Success Operation End

Displays information only for those volumes that match the specified last successful operation end time. This parameter is not supported on Infinite Volumes.

[-last-op-begin <Date>] - Last Operation Begin

Displays information only for those volumes that match the specified last operation begin time. This parameter is not supported on Infinite Volumes.

[-last-op-end <Date>] - Last Operation End

Displays information only for those volumes that match the specified last operation end time. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-last-op-error <text>] - Last Operation Error

Displays information only for those volumes that match the specified last operation error. This parameter is not supported on Infinite Volumes.

[-changelog-usage <percent_no_limit>] - Changelog Usage

Displays information only for those volumes that match the specified change log usage. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-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. This parameter is not supported on Infinite Volumes.

[-logical-data-percent <percent_no_limit>] - Logical Data Percent

Displays information only for those volumes that match the specified logical data percentage. This parameter is not supported on Infinite Volumes.

[-queued-job <text>] - Queued Job

Displays information only for those volumes that match the specified number of queued jobs. This parameter is not supported on Infinite Volumes.

[-stale-fingerprint-percentage <integer>] - Stale Fingerprint Percentage

Displays information only for those volumes that match the specified stale fingerprint percentage. This parameter is not supported on Infinite Volumes.

[-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.

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

Displays information only for those volumes that either are or are not constituents of an Infinite Volume, 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.

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 00:00:14
                                Type: Regular
                            Schedule: sun-sat@0
               Efficiency Policy Name: -
               Blocks Skipped Sharing: 0
                 Last Operation State: Success
        Last Success Operation Begin: Mon Nov 15 20:13:26 UTC 2010
          Last Success Operation End: Mon Nov 15 20:13:26 UTC 2010
                Last Operation Begin: Mon Nov 15 20:13:26 UTC 2010
                  Last Operation End: Mon Nov 15 20:13:26 UTC 2010
                 Last Operation Size: 0.00B
                Last Operation Error: -
                    Change Log Usage: 0%
                   Logical Data Size: 156KB
                  Logical Data Limit: 50.00TB
                Logical Data Percent: 0%
                           Queued Job: -
        Stale Fingerprint Percentage: 0
                         Compression: false
                  Inline Compression: false
                  Constituent Volume: false
                       Inline Dedupe: false
                     Data Compaction: false
   Cross Volume Inline Deduplication: false
ross Volume Background Deduplication: false
```

Related Links

- volume efficiency start

volume efficiency start

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.

This command is not supported on Infinite Volumes that are managed by storage services.

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) (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) (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) (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) (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) (privilege: advanced)

Compresses data blocks locked in a Snapshot copy. 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 (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
```

volume efficiency stat

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. This command is not supported on Infinite Volumes. 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.

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
```

volume efficiency stop

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. This command is not supported on Infinite Volumes that are managed by storage services.

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
```

volume efficiency undo

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 Data 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.

This command is not supported on Infinite Volumes that are managed by storage services.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

This specifies the Vserver on which the volume is located.

{ `-volume <volume name>` - Volume Name (privilege: advanced)

This specifies the volume name.

| `-path </vol/volume>` - Volume Path (privilege: advanced) }

This specifies the volume path.

`[-C, -compression <true>]` - Decompress Data in the Volume (privilege: advanced)

Undo the effects of compression. This requires efficiency to be disabled (by performing volume efficiency off).

`[-D, -dedupe <true>]` - Undo Block Sharing in the Volume (privilege: advanced)

Undo the effects of deduplication. This requires efficiency to be disabled (by performing volume efficiency off).

`[-i, -inode <integer>]` - Inode Number to Undo Sharing (privilege: advanced)

Remove the block sharings from a specified inode. This parameter is not supported on Infinite Volumes.

`[-t, -undo-type {all|wrong}]` - Selective Undo (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 (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 (privilege: advanced)

Undo the effects of data compaction.

[-A, -cross-volume-dedupe <true>] - Undo Cross Volume Deduplication (privilege: advanced)

Undo the effects of cross volume deduplication.

Examples

The following are examples of how to use efficiency undo.

To undo deduplication savings, but not compaction or compression 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 an SVM named vs1:

```
cluster1::> volume efficiency undo -vserver vs1 -volume vol1 -data
-compaction
```

Related Links

- volume efficiency off

volume efficiency policy create

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 <Efficiency policy type>] - 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 <Efficiency QoS policy>] - 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
```

volume efficiency policy delete

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
```

volume efficiency policy modify

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 <Efficiency policy type>]` - 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 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 <Efficiency QoS policy>] - 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

- job schedule show

volume efficiency policy show

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 Data 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 <Efficiency policy type>] - 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 <Efficiency QoS policy>] - 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.
```

volume encryption commands

volume encryption prepare-to-downgrade

Prepares volume encryption for releases earlier than Data ONTAP 9.1.0

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

Description

The `volume encryption prepare-to-downgrade` command prepares nodes to downgrade to a release earlier than ONTAP 9.1.0. Prior to disabling the features, the command checks whether the features are in use and if they are then, the command provides instructions for removing the use of the features.

Examples

The following example will disable the Volume Encryption features that are new to Data ONTAP *`9.1.0`* in the local cluster:

```
cluster1::> volume encryption prepare-to-downgrade
```

volume encryption conversion pause

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

volume encryption conversion resume

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

volume encryption conversion show

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 is in
                                              progress.
```

volume encryption conversion start

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

volume encryption rekey pause

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

volume encryption rekey resume

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

volume encryption rekey show

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 is in
                                              progress.
```

volume encryption rekey start

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

volume file commands

volume file compact-data

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 copy 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 copy 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 copy name in which the file will be compacted.

Examples

The following command applies the Adaptive Data Compaction feature to the Snapshot copy *snap1* of the file */file1* in volume *vol1* :

```
cluster1::> volume file compact-data -vserver vs1 -volume vol1 -file
/vol/vol1/file1 -snapshot snap1
```

volume file modify

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.

This command is not supported on FlexGroups or Infinite Volumes.

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”.

`[-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.

volume file privileged-delete

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.
```

volume file reservation

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. This command is not supported on Infinite Volumes. 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.
```

volume file show-disk-usage

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. This command is not supported on an Infinite Volume.

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
```

volume file show-filehandle

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 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)

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] }

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

[-vserver <vserver name>] - Vserver Managing Volume (privilege: advanced)

This specifies the Vserver where the file resides.

[-path <text>] - Path to File (privilege: advanced)

This specifies the path to the file.

Examples

The following example displays the file handle 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 file show-inode

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 copy; 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 copy 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 copy 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 copy 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 copy named mysnap. The Snapshot copy 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 copy named mysnap. The Snapshot copy 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.
```

volume file clone autodelete

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 a file or LUN clone. Newly created file and LUN clones are disabled for automatic deletion by default. This command is not supported on Infinite volumes.

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 or LUN 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 or LUN clone. Otherwise, specify the absolute path.

`-enable {true|false}` - Enable or Disable Autodelete

This parameter enables or disables the autodelete feature for the file or LUN clones in a specified volume if the clones are already added for automatic deletion. If you set the parameter to true, the specified file or LUN clones gets 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 or LUN, or a file or LUN clone. If `-enable` is *`false`* then specifying this parameter disables autodeletion on a file or LUN - or a file or LUN 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
```

volume file clone create

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 or a LUN. This command is not supported on Infinite Volumes. 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 or LUN clone created for auto deletion
- The option to overwrite an existing file or LUN clone

File or LUN clones create a duplicate copy of another file or LUN, 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 or LUN.

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 or LUN is going to be cloned.

`-source-path <text>` - Source Path

This specifies the path to the file or LUN to be cloned relative to the specified volume.

`-destination-path <text>` - Destination Path

This specifies the path for the newly-created cloned file or LUN relative to the specified volume. If the file or LUN clone to be created is a whole file or LUN, the destination file or LUN 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 copy to use as the source for the clone operation. If this value is not specified, the active filesystem will be used instead.

{ `[-r, -range [source start block>:<destination start block>:<block length],…]` - Block Range

This specifies the block range to be cloned. If the range is not specified, the entire file or LUN 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 copy respectively. If this parameter is used in the path provided by the `destination-path` parameter must refer to a file or LUN which already exists. If either the source or destination are a LUN then the block size is measured in 512-byte LBA blocks. If neither the source nor destination are a LUN 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>] - 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>] - Ignore streams

This parameter specifies whether streams should be ignored during cloning of files or LUNs. 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>] - Ignore locks

This parameter specifies whether byte-range locks and shared-mode locks on files or LUNs 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 or LUNs. The default value is false.

[-d, -overwrite-destination <true>] - Overwrite Destination

Specify this parameter to overwrite the destination file, if it exists. The default is to fail the command if the destination 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.

[-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 or LUN clones created for auto deletion. When set to true, the file or LUN clones get automatically deleted when the volume runs out of space. The default value is false.

[-bypass-throttle {true|false}] - Bypass Throttle Checks (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.

[-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 parameter 'volume', whereas parameter 'volume' specifies the volume on which source file resides. This is an optional argument applicable only for Metawafl volume where the source and destination volumes for the clone operation can be different. If this paramerter is not given the destination file will gets 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
```

volume file clone show-autodelete

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 or LUN clone. The command displays the following information about a file or LUN clone:

- Vserver Name
- Clone Path
- Whether auto deletion of file or LUN 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 or LUN clone belongs.

-clone-path <text> - Clone Path

This specifies the path of the file or LUN clone.

[-autodelete-enabled {true|false}] - Autodelete Enabled

If this parameter is true, the file or LUN 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 or LUN 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
```

volume file clone deletion add-extension

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. This command is not supported on Infinite Volumes.

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
```

volume file clone deletion modify

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. This is not applicable on Infinite volumes.

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
```

volume file clone deletion remove-extension

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. This command is not supported on Infinite Volumes.

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
```

volume file clone deletion show

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
```

volume file clone split load modify

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 or LUN 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
```

volume file clone split load show

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 or LUN 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 or LUN 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 or LUN 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
```

volume file fingerprint abort

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
```

volume file fingerprint dump

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
            Formatted Fingerprint End Time:Thu Apr 14 05:43:06 GMT 2016
                           Litigation Count:0
```

volume file fingerprint show

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.
```

volume file fingerprint start

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.
```

volume file retention show

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

Specifies the name of the Vserver which has the file.

[-file </vol/<volume name>/<file path>>] - File path

Specifies the absolute path of the file. The value begins with /vol/<volumename>.

[-retention-time <integer>] - Retention Time of the File

If this paramater is specified, the command returns the retention time of the file protected by SnapLock if its retention time in seconds since 1 January 1970 00:00:00 matches the specified value.

[-formatted-retention-time <text>] - Formatted Retention Period

If this paramater is specified, the command returns the retention time of the file protected by SnapLock if its expiry date matches the specified value. 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.

[-is-wraparound {true|false}] - Is Retention Time Wraparound

If this paramater is specified, the command returns the retention time of the file protected by SnapLock if it has a matching `-is-wraparound` value. 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.

Examples

The following example displays the retention time of the file */vol/nfs_sle/f12* :

```
cluster1::> volume file retention show -vserver vs0 -file /vol/nfs_sle/f12
Vserver : vs0
                                      Path : /vol/nfs_sle/f12
          Retention Time (Secs from Epoch) : 1439111404
                  Formatted Retention Time : Sun Aug  9 09:10:04 GMT 2015
              Is Retention Time Wraparound : false
```

volume flexgroup commands

volume flexgroup deploy

Deploy a FlexGroup on the cluster

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

Description

The `volume flexgroup deploy` command deploys a FlexGroup on the cluster.

The FlexGroup is created with 8 constituents on each node in the cluster. The constituents are split equally between the two largest aggregates on each node. If using two aggregates per node is not possible, all of the constituents are created on the largest aggregate on each node.

The `volume flexgroup deploy` command is only supported on clusters with 4 nodes or less. On clusters with more than 4 nodes, use the volume create command to create FlexGroups.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver on which the FlexGroup will be located.

`-size {<integer>[KB|MB|GB|TB|PB]}` - Size of the FlexGroup

This specifies the size of the FlexGroup. The size is specified as a number followed by a unit designation: k (kilobytes), m (megabytes), g (gigabytes), t (terabytes), or p (petabytes). 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 is 160 MB multiplied by the number of nodes in the cluster. The maximum size of the FlexGroup is limited by the platform maximum FlexVol size multiplied by 8 and multiplied by the number of nodes in the cluster. The size of the FlexGroup is also limited by the available space in the hosting aggregates. FlexGroups can be increased in size with the volume modify command and more constituents can be added with the volume expand command.

`[-volume <volume name>]` - Name of the FlexGroup to Create

This specifies the name of the FlexGroup. The name of the FlexGroup must start with an alphabetic character (a to z or A to Z) and must have 197 or fewer characters. The default value for the FlexGroup name is "fg".

`[-type {RW|DP}]` - Volume Type

This optionally specifies the FlexGroup'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.

[-space-guarantee {none|volume}] - Space Guarantee Style

This optionally specifies the space guarantee style for the FlexGroup. 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. Because CIFS does not handle out-of-space conditions, do not use the value *none* if the volume is accessible to CIFS clients. The default setting for volumes on All Flash FAS systems is *none* , otherwise the default setting is *volume* .

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

If this parameter is specified with *false* , the command runs in the background as a job. Otherwise, the command does not return until the operation is complete. The default value is *true*

Examples

The following example deploys a FlexGroup named "flexgroup" in the Vserver named "vs1.example.com":

```
cluster::> volume flexgroup deploy -size 400TB -vserver vs1.example.com
-volume flexgroup
```

Related Links

- volume create
- volume modify
- volume expand

volume flexgroup qtree-disable

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 (privilege: advanced)

This specifies the Vserver in which the FlexGroup is located.

-volume <volume name> - Volume Name (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
```

volume inode-upgrade commands

volume inode-upgrade prepare-to-downgrade

Prepare volume to downgrade to a release earlier than Data ONTAP 9.0.0

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

Description

The `volume inode-upgrade prepare-to-downgrade` command prepares volumes to downgrade to a release earlier than Data ONTAP 9.0.0. It is used when there are still volumes in the middile of the inode upgrade process when revert is issued.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

This specifies the node on which the command will run. Default is the local node.

Examples

The following example prepares volumes to revert to an earlier release.

```
cluster1::> volume inode-upgrade prepare-to-downgrade -node node1
```

volume inode-upgrade resume

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 (privilege: advanced)

This specifies the Vserver on which the volume is located.

`-volume <volume name>` - Volume Name (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
```

volume inode-upgrade show

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] }

If this parameter is specified, the command displays information about all entries.

[-vserver <vserver name>] - Vserver (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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>h][<integer>m][<integer>s]>] - Remaining Upgrade Time (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 (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
```

volume move commands

volume move abort

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 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

- job stop

volume move modify

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 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.

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 (privilege: advanced)

Specifies the action to be taken for cutover. If the effective cluster version is Data 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 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 volume move trigger-cutover command as the signal to try the cutover.

[-cutover-window <integer>] - Specified Cutover Time Window (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.

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

- volume move show
- volume move trigger-cutover

volume move show

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 (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 (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 (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 (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_soft_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 (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 (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 (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 (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.

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.
```

volume move start

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 (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 (privilege: advanced)

Specifies the action to be taken for cutover. If the effective cluster version is Data 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 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 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 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 {snapshot-only|none|backup}] - 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 capacity tier when they become cold. FabricPool combines flash (performance tier) with an object store (external capacity tier) into a single aggregate. 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 Snapshot copies not associated with the active file system.
- none - Volume blocks will not be tiered to the capacity tier.
- backup - This policy transfers all user data blocks directly to the capacity tier. In the case of DP volumes, the tiering policy is retained as backup after the move. For non-DP volumes, the policy is switched back to snapshot-only.

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

- volume move trigger-cutover

volume move trigger-cutover

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

volume move recommend show

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] }

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

[-vserver <vserver name>] - Vserver Name (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 (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 (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 (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 (%) (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 (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 (%) (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%
```

volume move target-aggr show

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 {snapshot-only|none|backup}] - 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”.

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.
```

volume qtree commands

volume qtree create

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. This command is not supported on Infinite Volumes.

You can optionally specify the following attributes when creating a new qtree:

- Security style
- Opportunistic lock mode
- 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). The *unified* security style, which applies only to Infinite Volumes, cannot be applied to a qtree. 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.

[-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 vserver export-policy create command. If you do not specify a value for this parameter, it inherits the export policy of its containing volume.

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

- vserver export-policy create

volume qtree delete

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 job show and job watch-progress commands, respectively. This command is not supported on Infinite Volumes.

The automatically created qtree in the volume - qtree0, listed in CLI output as "" - cannot be deleted.

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 (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

- job show
- job watch-progress

volume qtree modify

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
- UNIX permissions
- Export policy

This command is not supported by Infinite Volumes.

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). The *unified* security style, which applies only to Infinite Volumes, cannot be applied to a qtree. 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.

[-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 vserver export-policy create command.

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

- vserver export-policy create

volume qtree oplocks

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. This command is not supported on Infinite Volumes.

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
```

volume qtree rename

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. This command is not supported on Infinite Volumes.

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
```

volume qtree security

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. This command is not supported on Infinite Volumes.

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). The *`unified`* security style, which applies only to Infinite Volumes, cannot be applied to a qtree. 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
```

volume qtree show

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. This command is not supported on Infinite Volumes.

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:

- 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. The *unified* security style, which applies only to Infinite Volumes, cannot be applied to a qtree.

[-oplock-mode {enable|disable}] - Oplock Mode

Selects information about the qtrees that have the specified oplock mode.

[-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.

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.
```

volume qtree statistics-reset

Reset qtree statistics in a volume

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

Description

This command resets qtree statistics for all qtrees in a volume. This command is not supported on Infinite Volumes.

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 qtrees whose statistics you want to reset.

Examples

The following example resets statistics for all qtrees on the volume named vol1 on the Vserver named vs0.

```
cluster1::> volume qtree statistics-reset -vserver vs0 -volume vol1
```

volume qtree statistics

Display qtree statistics

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

Description

This command displays NFS and CIFS operations statistics for qtrees. Note that qtree statistics are available only when the volume containing the qtree is online. This command is not supported on Infinite Volumes.

Statistics are cumulative values from the time the volume is brought online or when the statistics have been reset by using the "volume qtree statistics-reset " command.

The command output depends on the parameters specified with the command. If no parameters are specified, the command displays the following statistics information about all qtrees:

- Vserver name
- Volume name
- Qtree name
- NFS operations
- CIFS operations

Qtree statistics are not persistent. If you restart a node, if a storage takeover and giveback occurs, or if the containing volume is set to offline and then online, qtree statistics are set to zero.

Parameters

{ [-fields <fieldname>,…]

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

| [-internal] (privilege: advanced)

If this parameter is specified, the output will also include the internal operation statistics. Internal operation is any operation on the qtree that originated within Data ONTAP software.

| [-no-reset] (privilege: advanced)

If this parameter is specified, the output will display the NFS and CIFS op statistics since the time the volume was online.

| [-no-reset-internal] (privilege: advanced)

If this parameter is specified, the output will also include the internal op statistics since the time the volume was online.

| [-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 about the qtrees on the specified Vserver.

{ [-volume <volume name>] - Volume Name

If this parameter is specified, the command displays information about the qtrees on the specified volume.

[-qtree <qtree name>] - Qtree Name

If this parameter is specified, the command displays information about the specified qtree.

| [-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>.

Examples

The following example displays statistics information for all qtrees on the Vserver named vs0.

```
cluster1::> volume qtree statistics -vserver vs0
Vserver    Volume        Qtree        NFS Ops      CIFS Ops
---------- ------------- ------------ ------------ ----------
vs0        vol0          qtree1       10876        2678
vs0        vol1          qtree1a      16543        0
vs0        vol2          qtree2       0            0
vs0        vol2          qtree2a      0            0
4 entries were displayed.
```

The following example displays statistics information for qtrees on Vserver vs0 that have NFS ops more than 15000.

```
cluster1::> volume qtree statistics -vserver vs0 -nfs-ops >15000
Vserver    Volume        Qtree        NFS Ops      CIFS Ops
---------- ------------- ------------ ------------ ----------
vs0        vol1          qtree1a      16543        0
```

Related Links

- volume qtree statistics-reset

volume quota commands

volume quota modify

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 job show and job watch-progress commands. This command is not supported on Infinite Volumes.

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 volume quota on , volume quota off and 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 Data 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.

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

- job show
- job watch-progress
- volume quota on
- volume quota off

- volume quota resize

volume quota off

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. This command is not supported on Infinite Volumes. You can monitor the progress of the job by using the job show and 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

- job show
- job watch-progress

volume quota on

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. This command is not supported on Infinite Volumes. You can monitor the progress of the job by using the job show and 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

- job show
- job watch-progress

volume quota report

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 is not supported on Infinite Volumes.

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 <text>] - 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: -
```

volume quota resize

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. This command is not supported on Infinite Volumes. You can monitor the progress of the job by using the job show and job watch-progress commands.

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

- job show
- job watch-progress

volume quota show

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. This command is not supported on Infinite 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 FlexGroups, 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 Data 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 (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 (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 (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:"").
```

volume quota policy copy

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. This command is not supported on Infinite Volumes. 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
```

volume quota policy create

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. This command is not supported on Infinite Volumes.

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
```

volume quota policy delete

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. This command is not supported on Infinite Volumes. 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
```

volume quota policy rename

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. This command is not supported on Infinite Volumes. 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
```

volume quota policy show

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. This command is not supported on Infinite Volumes. 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.
```

volume quota policy rule create

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. This command is not supported on Infinite Volumes. 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

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 "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 user -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

- volume quota modify

volume quota policy rule delete

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. This command is not supported on Infinite Volumes. 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

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 "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

- volume quota modify

volume quota policy rule modify

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

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 "volume quota modify " command.

```
This command is not supported on Infinite Volumes.
```

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

- volume quota modify

volume quota policy rule show

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

This command is not supported on Infinite 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 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
```

volume quota policy rule count show

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.
```

volume reallocation commands

volume reallocation measure

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, 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, file, or volume layout in the log file. To view previous measure-only reallocation checks, use the volume reallocation show command.

> This command is not supported for FlexGroups, FlexGroup Constituents, Infinite Volumes, or Infinite Volume constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, 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, 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

- volume reallocation show

volume reallocation off

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.

This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Examples

```
cluster1::> volume reallocation off
```

Disables all reallocation jobs globally in a cluster.

volume reallocation on

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.

This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Examples

```
cluster1::> volume reallocation on
```

Globally enables all reallocation jobs on a cluster.

volume reallocation quiesce

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 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.

This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver.

`-path <text>` - Path

Specifies the file path of the LUN, 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

- volume reallocation restart

volume reallocation restart

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.

This command is not supported on FlexGroups, FlexGroup constituents, or Infinite Volumes.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver.

`-path <text>` - Path

Specifies the file path of the LUN, 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.

volume reallocation schedule

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 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 volume reallocation start command.

> This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver.

`-path <text>` - Path

Specifies the path of the reallocation for a LUN, 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.

> 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

- volume reallocation start

volume reallocation show

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|Reschedule|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 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

- job schedule show

volume reallocation start

Start reallocate job

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

Description

Begins a reallocation scan on a LUN, 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, file, or volume.

You can define the reallocation scan job so that it runs at a specific interval, or you can use the volume reallocation schedule command to schedule reallocation jobs.

This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Parameters

`-vserver <vserver name>` - Vserver

Specifies the Vserver.

`-path <text>` - Path

Specifies the path of the reallocation for a LUN, 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, 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

- volume reallocation schedule

volume reallocation stop

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, file, or volume. This command stops and deletes in-progress, scheduled, and quiesced scans.

This command is not supported for FlexGroups, FlexGroup constituents, Infinite Volumes, or Infinite Volume constituents.

Parameters

-vserver <vserver name> - Vserver

Specifies the Vserver.

-path <text> - Path

Specifies the path of the reallocation for a LUN, file, or volume.

Examples

```
cluster1::> volume reallocation stop /vol/vol2
1 entry was deleted.
```

Stops and deletes one reallocation scan on volume vol2.

volume schedule-style commands

volume schedule-style prepare-to-downgrade

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
```

volume snaplock commands

volume snaplock modify

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}] - 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.

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> hours|days} | 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 days and hours. For example, a value of 2hours represents an autocommit period of

2 hours. The minimum allowed autocommit period is 2 hours and the maximum allowed autocommit period is 7 days.

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 Snapshot copies.

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 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

- volume file privileged-delete

volume snaplock prepare-to-downgrade

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
```

volume snaplock show

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

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}] - 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> hours|days} | 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.

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

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
```

volume snapshot commands

volume snapshot compute-reclaimable

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 Snapshot copies 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 Snapshot copies per query.

Snapshot copies must be specified as a comma-separated list with no spaces after the commas.

This command is not supported on Infinite 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 for which reclaimable space is to be calculated.

-snapshots <snapshot name>,… - List of Snapshots

This specifies one or more than one Snapshot copies that are to be considered for deletion. If you list more than one Snapshot copy, 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 copy 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
```

volume snapshot create

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 copy 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 copy is to be created.

-snapshot <snapshot name> - Snapshot

This specifies the name of the Snapshot copy that is to be created.

[-comment <text>] - Comment

This optionally specifies a comment for the Snapshot copy.

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

If you use this option and select *false* , the Snapshot copy creation process runs in the background. If you use this option and select *true* , the Snapshot copy creation process runs in the foreground. This option applies only to Infinite Volumes, and is ignored for other volumes. The default is *true* .

[-snapmirror-label <text>] - Label for SnapMirror Operations

If you specify this option, the Snapshot copy is created with the SnapMirror Label that you specify. If this option is not specified, the Snapshot copy is created with no SnapMirror Label. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-expiry-time <MM/DD/YYYY HH:MM:SS>] - Expiry Time

If you specify this option, the Snapshot copy is created with the expiry time that you specify. The expiry time indicates the time at which the Snapshot copy becomes eligible for deletion.

Examples

The following example creates a Snapshot copy named vol3_snap on a volume named vol3 on a Vserver named vs0. The Snapshot copy 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
```

volume snapshot delete

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 copy from a specified volume.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is saved.

-volume <volume name> - Volume

This specifies the volume from which a Snapshot copy is to be deleted.

-snapshot <snapshot name> - Snapshot

This specifies the Snapshot copy 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. This option applies only to Infinite Volumes, and is ignored for other volumes. The default is *true* .

[-force <true>] - Force Delete (privilege: advanced)

If you use this switch, the Snapshot copy 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 (privilege: advanced)

If you use this switch, the command ignores other processes that might be accessing the Snapshot copy. 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 copy named vol3_daily from a volume named vol3 on a Vserver named vs0:

```
cluster1::> volume snapshot delete -vserver vs0 -volume vol3 -snapshot
vol3_daily
```

volume snapshot modify-snaplock-expiry-time

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 copy.

Parameters

-vserver <vserver name> - Vserver

This specifies the Vserver that contains the volume on which the Snapshot copy is located.

-volume <volume name> - Volume

This specifies the volume where a Snapshot copy is to be located.

-snapshot <text> - Snapshot

This specifies the name of the Snapshot copy 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 copy locked by SnapLock.

If this option value is *infinite* , then a retention period that never expires is applied to the Snapshot copy.

Examples

The following example extends the retention period of a Snapshot copy *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 copy *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
```

volume snapshot modify

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 copy.

This command is not supported on Infinite Volumes.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is saved.

`-volume <volume name>` - Volume

This specifies the volume whose Snapshot copy is to be modified.

`-snapshot <snapshot name>` - Snapshot

This specifies the Snapshot copy whose text comment is to be modified.

[-comment <text>] - Comment

This specifies the new comment for the Snapshot copy.

[-snapmirror-label <text>] - Label for SnapMirror Operations

This specifies the SnapMirror Label for the Snapshot copy. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies 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 copy. The expiry time indicates the time at which the Snapshot copy 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 copy 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"
```

volume snapshot partial-restore-file

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 copy. This command is intended to be used to restore particular pieces of LUNs and NFS or CIFS container files that are used by a host to store multiple sources of data. For example, a host may 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 volume snapshot restore-file command to restore normal user-level files. The volume for the partial-restore should be online during this operation.

This command is not supported on Infinite Volumes.

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 copy is saved.

-s, -snapshot <snapshot name> - Snapshot Name

This specifies the Snapshot copy which contains the version of file from which a range of bytes is restored. The source file must be present in the Snapshot copy.

-path <text> - Filepath

This specifies the relative path to the file which is partially restored from the Snapshot copy. You should specify the `-volume` option so that the file is searched and restored from the Snapshot copy of the specified volume. If you do not specify the `-volume` then the file is searched and restored from the Snapshot copy of the root volume. The destination file 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 copy *`vol3_snap`* :

```
cluster1::> volume snapshot partial-restore-file -vserver vs0 -volume vol3
```

```
-snapshot vol3_snap -volume vol3 -path /foo.txt -start-byte 0 -byte-count
4096
```

Related Links

- volume snapshot restore-file

volume snapshot prepare-for-revert

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 Snapshot copies that have the format used by the current version of Data ONTAP. It will fail if any Snapshot copy polices are enabled, or if any Snapshot copies have an owner.

Snapshot policies must be disabled prior to running this command.

Parameters

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

The name of the node.

Examples

The following example prepares the Snapshot copies for revert.

```
cluster1::*> volume snapshot prepare-for-revert -node node1
```

volume snapshot rename

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 copy.

> You cannot rename a Snapshot copy that is created as a reference copy during the execution of the `volume move` command.

This command is not supported on Infinite Volumes.

Parameters

`-vserver <vserver name>` - Vserver

This specifies the Vserver that contains the volume on which the specified Snapshot copy is to be renamed

`-volume <volume name>` - Volume

This specifies the volume that contains the Snapshot copy to be renamed.

`-snapshot <snapshot name>` - Snapshot

This specifies the Snapshot copy that is to be renamed.

`-new-name <snapshot name>` - Snapshot New Name

This specifies the new name for the Snapshot copy.

`[-force <true>]` - Force Rename (privilege: advanced)

If this parameter is specified, the Snapshot copy rename operation is allowed on application tagged volumes. Otherwise, the operation is disallowed on application tagged volumes.

Examples

The following example renames a Snapshot copy named vol3_snap on a volume named vol3 and a Vserver named vs0. The Snapshot copy is renamed to vol3_snap_archive.

```
cluster1::> volume snapshot rename -vserver vs0 -volume vol3
-snapshot vol3_snap -new-name vol3_snap_archive
```

volume snapshot restore-file

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 copy. 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 copy. If the destination file for the restore operation exists, then it is overwritten by the version from the Snapshot copy. This operation is used to restore normal user-level files and LUNs. 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.

This command is not supported on Infinite Volumes.

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 copy.

`-s, -snapshot <snapshot name>` - Snapshot Name

This specifies the Snapshot copy from which the file is restored.

`-path <text>` - Filepath

This specifies the relative path to the file which is restored from the Snapshot copy. You should specify the `-volume` option so that the file is searched and restored from the Snapshot copy of the specified volume. If you do not specify the `-volume` then the file is searched and restored from the Snapshot copy 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 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 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 copy *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
```

volume snapshot restore

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 copy to be the read-write parent volume for the volume family. This replaces the current working copy of the volume with the Snapshot copy that results in a loss of all changes made since the Snapshot copy was created.

> You should manually update all the SnapMirror relationships of a volume immediately after you restore its Snapshot copy. Not doing so can result in unusable SnapMirror relationships that must be deleted and re-created.

Before running this command on an Infinite Volume, unmount the volume. Any namespace mirror constituents present in the system are resyncronized to the restored Snapshot copy.

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 copy to be restored is saved.

-volume <volume name> - Volume

This specifies the parent read-write volume whose Snapshot copy is to be restored to take its place.

-snapshot <snapshot name> - Snapshot

This specifies the Snapshot copy that is to be restored to be the read-write parent volume.

[-force <true>] - Force Restore

If you use this parameter, the Snapshot copy is restored even if the volume has one or more newer Snapshot copies which are currently used as reference Snapshot copy 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 snapmirror resync command if a common Snapshot copy is found between the source and destination volume. If there is no common Snapshot copy 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 copy restore needs to be non-disruptive to clients due to LUN 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 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 offline before proceeding.

Examples

The following example restores a Snapshot copy 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

- snapmirror resync

volume snapshot show-delta

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 Snapshot copies or a Snapshot copy 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 Snapshot copies in seconds.

Queries that use query operators (`*`, `|`, etc.) are disabled for this command to avoid performance degradation for client requests.

This command is not supported on Infinite 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 for which the delta is to be calculated.

-snapshot1 <snapshot name> - First Snapshot Name

This specifies the first Snapshot copy for the comparison.

[-snapshot2 <snapshot name>] - Second Snapshot Name

This specifies the second Snapshot copy 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 copies:

```
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 Snapshot copies: 1s.
```

volume snapshot show

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 Snapshot copies. 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 copies:

- Vserver name
- Volume name
- Snapshot copy 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 Snapshot copies, the detailed list view provides the following additional information:

- Creation time

- Snapshot busy
- List of the Snapshot copy's owners
- Comment associated with the Snapshot copy
- SnapMirror Label associated with the Snapshot copy
- 7-Mode Snapshot
- Constituent Snapshot
- Expiry Time
- SnapLock Expiry Time

At the advanced or higher privilege level the detailed view provides the following additional information:

- Snapshot copy's Dataset ID
- Snapshot copy's master Dataset ID
- Number of consistency points in the Snapshot copy
- Internal status of the Snapshot copy
- 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
- Status of FlexGroup Qtree support in the Snapshot copy

For Snapshot copies 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 Snapshot copies whose parent volume is
a FlexGroup Constituent.
```

At the admin and advanced privilege level, Snapshot copies whose parent volume is a FlexGroup Constituent are not displayed by default. To display these, run the commmand and set the `is-constituent` to *`true`* . At the diagnostic or higher privilege level, all Snapshot copies are displayed by default.

The list view is automatically enabled if a single Snapshot copy 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 Snapshot copies:

- Vserver name
- Volume name
- Snapshot copy 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 Snapshot copies of the load-sharing volumes, run the command with the `-volume-type LS` parameter. If you specify multiple filtering parameters, only those Snapshot copies 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 Snapshot copies located only on the specified Vserver will be displayed.

[-volume <volume name>] - Volume

If you use this parameter only Snapshot copies located on the specified volume will be displayed.

[-snapshot <snapshot name>] - Snapshot

If you use this parameter only Snapshot copies matching the specified name will be displayed.

[-dsid <integer>] - Snapshot Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copy that has the specified data set ID.

[-msid <integer>] - Snapshot Master Data Set ID (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copy 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 Snapshot copies that match the specified creation time.

[-busy {true|false}] - Snapshot Busy

If this parameter is specified, the command displays information only about the Snapshot copies that have the specified busy status.

[-owners <text>,…] - List of Owners

If this parameter is specified, the command displays information only about the Snapshot copies 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 Snapshot copies that have the specified size. The size is specified as a character specifying the unit of measurement followed by 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 Snapshot copies 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 Snapshot copies 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 (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copies 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 Snapshot copies that have the specified comment text. You may also specify an inequality such as *!"-"* as input.

[-fs-version <text>] - File System Version (privilege: advanced)

If you use this parameter the only Snapshot copies 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 Snapshot copies that will be impacted by the upgrade process.

[-is-7-mode {true|false}] - 7-Mode Snapshot

If you use this parameter only those Snapshot copies which have the specified value are shown. This value is *true* for the Snapshot copies 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 Snapshot copies are still in the 7-mode configuration.

[-snapmirror-label <text>] - Label for SnapMirror Operations

If you use this parameter, only those Snapshot copies that have the specified SnapMirror Label value are shown.

[-state {valid|invalid|partial}] - Snapshot State

If you use this parameter only those Snapshot copies which have the specified state will be shown.

[-is-constituent {true|false}] - Constituent Snapshot

If you use this parameter, only those Snapshot copies whose parent volume is a constituent volume of a FlexGroup or an Infinite Volume will be shown.

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

If you use this parameter only those Snapshot copies that are located on the specified storage system are shown.

[-inofile-version <integer>] - Snapshot Inofile Version (privilege: advanced)

If this parameter is specified, the command displays information only about the Snapshot copies 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 Snapshot copies that have the specified expiry time are shown.

[-compression-type {none|secondary|adaptive}] - Compression Type (privilege: advanced)

If you use this parameter only those Snapshot copies 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 Snapshot copies that have the specified snaplock expiry time are shown.

[-application-io-size {default|8K|16K|32K}] - Application IO Size (privilege: advanced)

If you use this parameter only those Snapshot copies that have the specified application IO size are shown.

[-is-flexgroup-qtree-enabled {true|false}] - Is FlexGroup Qtree Support Enabled (privilege: advanced)

If you use this paramater only those Snapshot copies that have the specified FlexGroup Qtree support status are shown.

Examples

The following example displays default information about all Snapshot copies 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 Snapshot copies 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 copy, 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: -
```

volume snapshot autodelete modify

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 or file clone autodelete policy settings. Based on the defined policy, automatic deletion of Snapshot copies and LUN or file clones is triggered. Automatic deletion of Snapshot copies and LUN or file clones is useful when you want to automatically reclaim space consumed by the Snapshot copies and LUN or file clones from the volume when it is low in available space. LUN or file clone autodelete follows Snapshot copy autodelete. This command works only on a read-write parent volume. You cannot setup automatic Snapshot copy deletion and automatic LUN or file clone deletion for Infinite Volumes and 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 Snapshot copies and LUN or file clones is enabled or disabled. If set to *true* , automatic deletion of Snapshot copies and LUN or file clones is enabled. If set to *false* , automatic deletion of Snapshot copies and LUN or file clones is disabled.

[-commitment {try|disrupt|destroy}] - Commitment

This option specifies which Snapshot copies and LUN or file clones can be automatically deleted to reclaim back space. + When set to *try* , the Snapshot copies which are not locked by any application and the LUN or file clones which are not configured as preserved are deleted. + When set to *disrupt* , the Snapshot copies which are not locked by data backing functionalities (such as volume clones, LUN clones and file clones) and LUN or file clones which are not configured as preserved are deleted. In the *disrupt* mode, the Snapshot copies locked by data protection utilities such as Snapmirror and Volume Move can be deleted. If such a locked Snapshot copy is deleted during the data transfer, the transfer is aborted. + When set to *destroy* , the Snapshot copies locked by the data backing functionalities are deleted. In addition, all the LUN or file clones in the volume are deleted.

[-defer-delete {scheduled|user_created|prefix|none}] - Defer Delete

This option determines the order in which Snapshot copies can be deleted. + Possible values are as follows:

- When set to *scheduled* , scheduled Snapshot copies are the last to be deleted.
- When set to *user_created* , user Snapshot copies are the last to be deleted.
- When set to *prefix* , Snapshot copies 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 or file clones.
```

[-delete-order {newest_first|oldest_first}] - Delete Order

This option specifies if the oldest Snapshot copy and the oldest LUN or file clone or the newest Snapshot copy and the newest LUN 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 or file clones.

[-target-free-space <percent>] - Target Free Space

This option specifies the free space percentage at which the automatic deletion of Snapshot copies and LUN or file clones must stop. Depending on the `-trigger` Snapshot copies and LUN 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 Snapshot copies and LUN or file clones. + Setting this option to *volume* triggers automatic deletion of Snapshot copies and LUN or file clones when the volume reaches threshold capacity and the volume space reserved for Snapshot copies is exceeded. + Setting the option to *snap_reserve* triggers automatic deletion of Snapshot copies and LUN or file clones when the space reserved for Snapshot copies reaches threshold capacity. + Setting the option to (DEPRECATED)-*space_reserve* triggers automatic deletion of Snapshot copies when reserved space

in the volume reaches threshold capacity and the volume space reserved for Snapshot copies is exceeded. +

> 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 copy 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. + If you specify *vol_clone* , the cloned volume backed by the Snapshot copy is deleted. + If you specify *lun_clone* , and the LUN is in the process of being cloned when autodelete is triggered, the cloning operation is aborted. Any access to this LUN 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 copy 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 copy 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 copy 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 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 or file clone autodelete for volume *vol3* which is part of the Vserve *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
```

volume snapshot autodelete show

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 Infinite 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
```

volume snapshot policy add-schedule

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 job schedule cron create or 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 Snapshot copies that can be taken by the specified schedule. The total count of all the Snapshot copies to be retained for the policy cannot be more than 255.

[-prefix <text>] - Snapshot Copy Name Prefix for Schedule

This option specifies the prefix with which Snapshot copies 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 copy when it is created for the added schedule. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

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 Snapshot copies.

```
cluster1::> volume snapshot policy add-schedule -vserver vs0 -policy
snappolicy_nightly -schedule midnight -count 5
```

Related Links

- job schedule cron create
- job schedule interval create

volume snapshot policy create

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 Snapshot copies per schedule. You can create a schedule by using the job schedule cron create or job schedule interval create commands. When applied to a volume, the Snapshot policy specifies the schedule on which Snapshot copies are taken and the maximum number of Snapshot copies that each schedule can take. The total count of all the Snapshot copies to be retained for the policy cannot be more than 255.

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 Snapshot copies 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 Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-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 Snapshot copies 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 Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination.

[-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 Snapshot copies 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.

[-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 Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies 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 Snapshot copies 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.

[-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 Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies 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 Snapshot copies 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.

`[-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 Snapshot copies created for that schedule have the SnapMirror Label assigned to them. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies 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 Snapshot copies.

```
cluster1::> volume snapshot policy create -vserver vs0 -policy
snappolicy_4hrs
-schedule1 4hrs -count1 5 -prefix1 every_4_hour
```

Related Links

- job schedule cron create
- job schedule interval create

volume snapshot policy delete

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
```

volume snapshot policy modify-schedule

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 Snapshot copies 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 modifed.

`-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 Snapshot copies that can be taken by the specified schedule. The total count of all the Snapshot copies to be retained for the policy cannot be more than 255.

`[-newsnapmirror-label <text>]` - Label for SnapMirror Operations

This specifies the SnapMirror Label identified with a Snapshot copy when it is created for the specified schedule. The SnapMirror Label is used by the Vaulting subsystem when you back up Snapshot copies to the Vault Destination. If an empty label ("") is specified, the existing label will be deleted.

Examples

The following example changes the maximum number of Snapshot copies 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
```

volume snapshot policy modify

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"
```

volume snapshot policy remove-schedule

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
```

volume snapshot policy show

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 Snapshot copies associated with each schedule
- Snapshot copy 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] (privilege: advanced)

If this parameter is specified, the command displays Snapshot policies that are not supported in Data ONTAP 8.2. The total Snapshot copy 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 Snapshot copies 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 Snapshot copies.

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.
```

volume transition-convert-dir commands

volume transition-convert-dir show

Display 7-Mode directories being converted

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

Description

The `volume transition-convert-dir show` command displays information about ongoing directory copy conversion 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] }

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

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

Displays summary information about the ongoing copy conversions of directories for the volumes in the specified Vserver.

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

Displays summary information about the ongoing copy conversions of directories that are occurring on the specified volume.

[-path <text>] - Directory Being Converted (privilege: advanced)

Displays summary information for the ongoing copy conversions of directories that have the specified directory path to convert.

[-job-id <integer>] - Convert Job ID (privilege: advanced)

Displays summary information for the ongoing copy conversions of directories that have the specified job ID.

[-state {Initial|Queued|Running|Waiting|Pausing|Paused|Quitting|Success|Failure|Reschedule|Error|Quit|Dead|Unknown|Restart|Dormant}] - Operation State (privilege: advanced)

Displays summary information for the copy conversions of directories that have the specified job state.

[-bytes-total <integer>] - Bytes Total (privilege: advanced)

Displays summary information for copy conversions which have the estimated number of bytes of directory content to convert.

[-bytes-completed <integer>] - Bytes Completed (privilege: advanced)

Displays summary information for copy conversions which have the estimated number of bytes of directory content that have completed conversion. The value of this field will be updated approximately once per minute.

Examples

The following example illustrates how to show directory conversions for a volume:

```
cluster1::*> volume transition-convert-dir show -volume vol1
Vserver    Volume    Convert Job ID   Directory-path   State
---------  --------- ---------------- ---------------- ------------
vs0        vol1                  151 /data/large_dir       Running
```

volume transition-convert-dir start

Start converting a 7-Mode directory to Cluster-mode

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

Description

The `volume transition-convert-dir start` command moves the directory entries in an existing directory to a new temporary directory and then replaces the existing directory with the temporary directory. This command only has a use for directories that were created in a non-Unicode format on a 7-Mode storage system and then transitioned to clustered Data ONTAP by using a SnapMirror relationship of type TDP. This command converts the directories to the Unicode format in a way that is less likely to disrupt the operation of the Data ONTAP systems than the existing directory conversion mechanisms. The temporary directory is visible from clients. Attempting to manipulate the directory being copied or the temporary directory might result in expected side-effects and should be avoided.

Parameters

-vserver <vserver name> - Vserver Name (privilege: advanced)

Specifies the Vserver on which the volume is located.

-volume <volume name> - Volume Name (privilege: advanced)

Specifies the volume in which the directory to be converted is located.

-path <file path> - Directory Path (privilege: advanced)

Specifies the path to the directory to be converted from the root of the volume specified with the `-volume` parameter. The root directory of a volume might not be converted using this command. Also, the path must not have a symbolic link as the last component in the path.

Examples

The following example shows how to start a 7-mode directory conversion for a given path in a volume:

```
cluster1::*> volume transition-convert-dir start -vserver vs0  -volume
vol1 -path /data/large_dir
```

vserver commands

vserver add-aggregates

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
```

vserver add-protocols

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} - Protocols

This parameter specifies the list of protocols to be allowed to run on the Vserver. Possible values include *nfs* , *cifs* , *fcp* , and *iscsi* , and *ndmp* .

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
```

vserver context

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 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

- exit

vserver create

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.

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 *`svm_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. The *`unified`* security style, which applies only to Infinite Volumes, cannot be applied to a Vserver's 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 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 optionally specifies a comment for the Vserver.

[-quota-policy <text>] - Quota Policy

This optionally specifies a quota policy for the Vserver. This parameter is not supported on a Vserver with Infinite Volume.

[-is-repository {true|false}] - Is Vserver with Infinite Volume

This parameter is deprecated and may be removed in a future release of Data ONTAP. If you are using Infinite Volumes it is recommended that you do not upgrade the cluster to a release that is later than Data ONTAP 9.3.0.

This specifies that the Vserver will contain an Infinite Volume.

[-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.

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 *svm_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     svm_root   aggr1     unix
```

Related Links

- volume create
- volume modify

vserver delete

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 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

- snapmirror break

vserver modify

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*. This field is not modifiable on a Vserver with Infinite Volume.

`[-snapshot-policy <snapshot policy>]` - Snapshot Policy

This optional parameter specifies the Snapshot policy for a Vserver being modified.

`[-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. This parameter is not supported on a Vserver with Infinite Volume.

`[-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 <unsigned32_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. This value is not modifiable on a Vserver with Infinite Volume.

[-admin-state {running|stopped|starting|stopping}] - Vserver Admin State (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}] - 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* , and *ndmp* . Possible values for a Vserver with Infinite Volume include *nfs* and *cifs* .

[-disallowed-protocols {nfs|cifs|fcp|iscsi|ndmp}] - 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* , and *ndmp* . Only the protocols configured for Vservers with Infinite Volume can be disallowed.

[-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". This parameter is not supported on a Vserver with Infinite Volume.

[-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.

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

This optionally specifies whether the Vserver modify operation can be executed in the background. If nothing is specified, by default the Vserver modify operation is executed in the foreground.

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
```

vserver prepare-for-revert

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 Data 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
```

vserver remove-aggregates

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
```

vserver remove-protocols

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}` - Protocols

This parameter specifies the list of protocols to be removed. on the Vserver. Possible values include *nfs*, *cifs*, *fcp*, *iscsi*, and *ndmp*.

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
```

vserver rename

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.

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 backgound. 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.
```

vserver restamp-msid

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 snapmirror resync must be run after this command completes.

Parameters

`-vserver <vserver name>` - Vserver name (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. (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

- snapmirror resync

vserver show-aggregates

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.
```

vserver show-protocols

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}] - 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
```

vserver show

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), or *`unified`* (Infinite Volumes 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`* , *`ndmp`* - detailed view only)
- Disallowed Protocols (*`nfs`* , *`cifs`* , *`fcp`* , *`iscsi`* , *`ndmp`* - detailed view only)

- Whether the Vserver is a Vserver with Infinite Volume (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. The *unified* security style, which applies only to Infinite Volumes, cannot be applied to a Vserver's root volume.

[-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.

[-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 <unsigned32_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}] - 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}] - 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}] - 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 specified, the command displays information only about the Vservers which have the specified is-repository value. This will be true for Vservers with Infinite Volume.

[-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.

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.
```

vserver start

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.

> 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
```

vserver stop

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.

> 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
```

vserver unlock

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 (privilege: advanced)

The name of the Vserver that has to be unlocked.

[-force <true>] - Force Unlock (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

vserver active-directory create

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.

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. Data 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.

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:
```

vserver active-directory delete

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:
```

vserver active-directory modify

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:
```

vserver active-directory password-change

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
```

vserver active-directory password-reset

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:
```

vserver active-directory show

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.

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
```

vserver audit commands

vserver audit create

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.

This command is not supported on a Vserver with Infinite Volume.

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}]` - 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 promt using 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 {<integer>[KB|MB|GB|TB|PB]}]` - 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.

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

- vserver audit modify

vserver audit delete

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 (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
```

vserver audit disable

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
```

vserver audit enable

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.

Events on FlexGroup volumes are not emitted to the audit log.

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 (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
```

vserver audit modify

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}]` - 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 {<integer>[KB|MB|GB|TB|PB]}]` - 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.

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
```

vserver audit prepare-to-downgrade

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 (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
```

vserver audit rotate-log

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
```

vserver audit show

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}] - 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 {<integer>[KB|MB|GB|TB|PB]}] - 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.

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
```

vserver check commands

vserver check lif-multitenancy run

Run check for LIF multitenancy

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

Description

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 (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.
```

vserver check lif-multitenancy show-results

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

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.
```

vserver check lif-multitenancy show

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

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
```

vserver cifs commands

vserver cifs add-netbios-aliases

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
```

vserver cifs create

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.

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.

`[-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 Data ONTAP cannot determine an appropriate site.

| `-workgroup <NetBIOS>` - Workgroup Name }

This parameter specifies the name of the workgroup (up to 15 characters).

`[-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 48 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:
```

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
```

vserver cifs delete

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.

Examples

The following example deletes the CIFS server from a Vserver named vs1:

```
cluster1::> vserver cifs delete -vserver vs1
```

vserver cifs modify

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 Data 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 Data 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).

`[-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 Data 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 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::>
```

vserver cifs nbtstat

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.

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.
```

vserver cifs prepare-to-downgrade

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 Data ONTAP based on the input parameter disable-feature-set.

Parameters

`-disable-feature-set <downgrade version>` - Data ONTAP Version (privilege: advanced)

This parameter specifies the Data 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 Data 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 Data 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 Data 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
```

vserver cifs remove-netbios-aliases

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: -
```

vserver cifs repair-modify

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 (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::*>
```

vserver cifs show

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
```

vserver cifs start

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 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

- vserver cifs create

vserver cifs stop

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.

> 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
```

vserver cifs branchcache create

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.

The `vserver cifs branchcache create` command is not supported for Vservers with Infinite Volume.

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 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

- vserver cifs share modify

vserver cifs branchcache delete

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.

The `vserver cifs branchcache delete` command is not supported for Vservers with Infinite Volume.

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
```

vserver cifs branchcache hash-create

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.

The `vserver cifs branchcache hash-create` command is not supported for Vservers with Infinite Volume.

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
```

vserver cifs branchcache hash-flush

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.

The `vserver cifs branchcache hash-flush` command is not supported for Vservers with Infinite Volume.

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
```

vserver cifs branchcache modify

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.

The `vserver cifs branchcache modify` command is not supported for Vservers with Infinite Volume.

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 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

- vserver cifs share modify

vserver cifs branchcache show

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
```

vserver cifs character-mapping create

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.

> Choose target characters in the "Private Use Area" of Unicode in the following range: U+E0000…U+F8FF.

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 FlexGroups or Vservers with Infinite Volume.

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.

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
```

vserver cifs character-mapping delete

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 FlexGroups or Vservers with Infinite Volume.

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
```

vserver cifs character-mapping modify

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

Choose target characters in the "Private Use Area" of Unicode in the following range: U+E0000…U+F8FF.

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 FlexGroups or Vservers with Infinite Volume.

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.

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
```

vserver cifs character-mapping show

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
```

vserver cifs connection show

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 (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
```

vserver cifs domain discovered-servers reset-servers

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::>
```

vserver cifs domain discovered-servers show

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|NIS}] - 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|winthreshold}] - 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.
```

vserver cifs domain discovered-servers discovery-mode modify

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.

Parameters

-vserver <vserver name> - Vserver (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 (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 servers in the domain.
- site - Discover the servers local to the site.
- 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
```

vserver cifs domain discovered-servers discovery-mode show

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 servers for the CIFS domains 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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you use this parameter, the command only displays discovery mode for the specified Vserver.

[-mode {all|site|none}] - Server Discovery Mode (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.
```

vserver cifs domain name-mapping-search add

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 Windows user to UNIX 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
```

vserver cifs domain name-mapping-search modify

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 Windows user to UNIX 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
```

vserver cifs domain name-mapping-search remove

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 Windows user to UNIX 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
```

vserver cifs domain name-mapping-search show

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
```

vserver cifs domain password change

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::>
```

vserver cifs domain password reset

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.

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::>
```

vserver cifs domain password schedule modify

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
```

vserver cifs domain password schedule show

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 accidently 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.
```

vserver cifs domain preferred-dc add

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.

> 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.

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
```

vserver cifs domain preferred-dc remove

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
```

vserver cifs domain preferred-dc show

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
```

vserver cifs domain trusts rediscover

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
```

vserver cifs domain trusts show

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
```

vserver cifs group-policy modify

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
```

vserver cifs group-policy show-applied

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.
```

vserver cifs group-policy show-defined

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
```

vserver cifs group-policy show

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
```

vserver cifs group-policy update

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
```

vserver cifs group-policy central-access-policy show-applied

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.
```

vserver cifs group-policy central-access-policy show-defined

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.
```

vserver cifs group-policy central-access-rule show-applied

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.
```

vserver cifs group-policy central-access-rule show-defined

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.
```

vserver cifs group-policy restricted-group show-applied

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 Data 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 {Local|Site|Domain|OrganizationalUnit|RSOP}] - 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
```

vserver cifs group-policy restricted-group show-defined

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 Data 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 defeined 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 {Local|Site|Domain|OrganizationalUnit|RSOP}] - 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
```

vserver cifs home-directory modify

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 (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
```

vserver cifs home-directory show-user

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
    ~GPO~rpuser1                             /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  ~GPO~rpuser1
Vserver  : vs1
    Username : GPO/rpuser1
ShareName                                Home Directory Path
    ------------------------------------
---------------------------------
    ~GPO~rpuser1                             /home/GPO/rpuser1
```

vserver cifs home-directory show

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 (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
```

vserver cifs home-directory search-path add

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
```

vserver cifs home-directory search-path remove

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
```

vserver cifs home-directory search-path reorder

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
```

vserver cifs home-directory search-path show

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
```

vserver cifs options modify

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.

Use an IPv4 address because WINS over IPv6 is not supported.

`[-smb1-enabled {true|false}]` - Enable SMB1 Protocol (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 true for Vservers with Infinite Volume. For other data Vservers, the default value is false.

`[-smb2-enabled {true|false}]` - Enable all SMB2 Protocols (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. This parameter is not supported for Vservers with Infinite Volume.

`[-smb3-enabled {true|false}]` - Enable SMB3 Protocol (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. This parameter is not supported for Vservers with Infinite Volume.

`[-smb31-enabled {true|false}]` - Enable SMB3.1 Protocol (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. This parameter is not supported for Vservers with Infinite Volume.

`[-max-mpx <integer>]` - Maximum Simultaneous Operations per TCP Connection (privilege: advanced)

This optional parameter specifies the maximum number of simultaneous operations the CIFS server reports it can process per TCP connection.

`[-shadowcopy-dir-depth <integer>]` - Maximum Depth of Directories to Shadow Copy (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 Vservers with Infinite Volume and workgroup CIFS servers. Directories and files within a FlexGroup will not be shadow copied because FlexGroups do not support shadow copy.

[-copy-offload-enabled {true|false}] - Enable Copy Offload Feature (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. This parameter is not supported for Vservers with Infinite Volume. Copy Offload has no effect on files in a FlexGroup because FlexGroups do not support Copy Offload.

[-is-copy-offload-direct-copy-enabled {true|false}] - Is Direct-copy Copy Offload Mechanism Enabled (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. This parameter is not supported for Vservers with Infinite Volume. Copy Offload has no effect on files in a FlexGroup because FlexGroups do not support Copy Offload. + 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) (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 Vservers with Infinite Volume and workgroup CIFS servers. Directories and files within a FlexGroup will not be shadow copied because FlexGroups do not support shadow copy.

[-is-referral-enabled {true|false}] - Refer Clients to More Optimal LIFs (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. This parameter is not supported for Vservers with Infinite Volume.

[-is-local-auth-enabled {true|false}] - Is Local User Authentication Enabled (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 (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 (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. This parameter is not supported for Vservers with Infinite Volume.

[-is-exportpolicy-enabled {true|false}] - Is Export Policies for CIFS Enabled (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 (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 (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 (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 Vservers with Infinite Volume and workgroup CIFS servers.

[-restrict-anonymous {no-restriction|no-enumeration|no-access}] - Restrictions for Anonymous User (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) (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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' (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 (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.

This parameter is deprecated and may be removed in a future release of Data 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 (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 (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 (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 (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 (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 (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-large-mtu-enabled {true|false}] - Is Large MTU Enabled (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 false. This parameter is not supported for Vservers with Infinite Volume.

[-is-netbios-over-tcp-enabled {true|false}] - Is NetBIOS over TCP (port 139) Enabled (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 (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 (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* .

Any values entered for this parameter is replaced with the existing values.

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
```

vserver cifs options show

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}] - Enable SMB1 Protocol (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.

[-smb2-enabled {true|false}] - Enable all SMB2 Protocols (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 (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 (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 (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 (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 (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 (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) (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 (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 (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 (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 (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 (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 (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 (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 (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 (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) (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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 (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' (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 (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.

This parameter is deprecated and may be removed in a future release of Data 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 (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 (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 (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 (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 (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 (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-large-mtu-enabled {true|false}] - Is Large MTU Enabled (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 (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 (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 (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.

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 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 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
```

vserver cifs security modify

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 is 5 minutes.

[-kerberos-ticket-age <integer>] - Kerberos Ticket Lifetime

This parameter specifies the Kerberos ticket lifetime in hours. The default 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 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 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 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 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 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-aes-encryption-enabled {true|false}] - 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 utilize these encryption types. If you do not specify this parameter, the default is *false* .

[-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 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. If you do not specify this parameter, the default 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. 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.

[-smb1-enabled-for-dc-connections {false|true|system-default}] - SMB1 Enabled for DC Connections

This parameter specifies whether SMB1 is enabled for use with connections to domain controllers. If you do not specify this parameter, the default 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 Data ONTAP that is running. For this release it is: SMB1 is enabled.

[-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. If you do not specify this parameter, the default 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 Data ONTAP that is running. For this release it is: SMB2 is enabled.

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:           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
```

vserver cifs security show

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}] - 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.

[-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}] - 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.

[-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.

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
```

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
```

vserver cifs session close

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.

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.
```

vserver cifs session show

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.

| [-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 <elapsed>] - Connected Time

If you specify this parameter, the command displays information about CIFS sessions that are established for the specified time duration.

[-idle-time <elapsed>] - 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.

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}] - 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.

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: -
```

vserver cifs session file close

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:

```
cluster1::> vserver cifs session file close -node node1 -vserver vs1
5 entries were acted on.
```

The following example closes all the opened CIFS files on all the nodes with the file-id 1:

```
cluster1::> vserver cifs session file close -node * -file-id 1
2 entries were acted on.
```

vserver cifs session file show

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.

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
```

vserver cifs share create

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 256 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, Data 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 Snapshot copies can be viewed and traversed by clients.
- changenotify - This property specifies that the share supports ChangeNotify requests. For shares on a Vserver with FlexVol volumes, this is a default initial property. For shares on a Vserver with Infinite Volume, the ChangeNotify property is not set by default, and setting it requires the advanced privilege level. When the ChangeNotify property is set for a share on a Vserver with Infinite Volume, change notifications are not sent for changes to file attributes and timestamps. If the path of the share is within a FlexGroup, change notifications are not sent because FlexGroups do not support ChangeNotify.
- attributecache - This property enables the file attribute caching on the CIFS share in order to provide faster access of attributes over SMB 1.0.

  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 FlexGroups, Vservers with Infinite Volume 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. This option is not supported for Vservers with Infinite Volume.
- 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}] - 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 (OBSOLETE)- This property enables clients to follow symlinks outside share boundaries.

* The *`enable`* , *`hide`* , and *`read-only`* parameters are deprecated and may be removed in a future release of Data ONTAP.

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>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.

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
```

vserver cifs share delete

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
```

vserver cifs share modify

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 256 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 FlexGroups.

[-symlink-properties {enable|hide|read-only|symlinks|symlinks-and-widelinks|disable}] - 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 (OBSOLETE)- This property enables clients to follow symlinks outside share boundaries.

The read_only setting does not apply to wide links.

* The *`enable`* , *`hide`* , and *`read-only`* parameters are deprecated and may be removed in a future release of Data ONTAP.

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>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.

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 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 not generated when a change occurs.

```
cluster1::> vserver cifs share modify -vserver vs1 -share-name SALES_SHARE
-path /sales -share-properties oplocks,browsable
```

```
The following example modifies a CIFS share named DOCUMENTS on a Vserver
named vs1. 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 modify -vserver vs1 -share-name DOCUMENTS
-share-properties branchcache,changenotify,oplocks
```

```
The following example modifies a CIFS share named DOCUMENTS on a Vserver
named vs1. 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 modify -vserver vs1 -share-name DOCUMENTS
-share-properties changenotify,oplocks -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
```

vserver cifs share show

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}] - 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>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
```

vserver cifs share access-control create

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
```

vserver cifs share access-control delete

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
```

vserver cifs share access-control modify

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
```

vserver cifs share access-control show

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.

[-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
```

vserver cifs share properties add

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 Snapshot copies can be viewed and traversed by clients.
- changenotify - This property specifies that the share supports ChangeNotify requests. For shares on a Vserver with FlexVol volumes, this is a default initial property. For shares on a Vserver with Infinite Volume, the ChangeNotify property is not set by default, and setting it requires the advanced privilege level. When the ChangeNotify property is set for a share on a Vserver with Infinite Volume, change notifications are not sent for changes to file attributes and timestamps. If the path of the share is within a FlexGroup, change notifications are not sent because FlexGroups do not support ChangeNotify.
- attributecache - This property enables the file attribute caching on the CIFS share in order to provide faster access of attributes over SMB 1.0.

  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 FlexGroups, Vservers with Infinite Volume 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. This option is not supported for Vservers with Infinite Volume.
- 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.

> 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
```

vserver cifs share properties remove

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 Snapshot copies can be viewed and traversed by clients.

- changenotify - This property specifies that the share supports ChangeNotify requests. For shares on a Vserver with FlexVol volumes, this is a default initial property. For shares on a Vserver with Infinite Volume, the ChangeNotify property is not set by default, and setting it requires the advanced privilege level. When the ChangeNotify property is set for a share on a Vserver with Infinite Volume, change notifications are not sent for changes to file attributes and timestamps. If the path of the share is within a FlexGroup, change notifications are not sent because FlexGroups do not support ChangeNotify.
- attributecache - This property enables the file attribute caching on the CIFS share in order to provide faster access of attributes over SMB 1.0.

  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 FlexGroups, Vservers with Infinite Volume 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. This option is not supported for Vservers with Infinite Volume.
- 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
```

vserver cifs share properties show

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 Snapshot copies 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.

  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 FlexGroups 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.
```

vserver cifs superuser create

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, Data ONTAP bypasses some of the security checks. This command is not supported for workgroup CIFS servers.

Parameters

`-vserver <vserver name>` - Vserver (privilege: advanced)

Vserver name.

`-domain <CIFS domain>` - Domain (privilege: advanced)

The domain name of accountname.

-accountname <CIFS account> - User (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
```

vserver cifs superuser delete

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, Data ONTAP bypasses some of the security checks.

Parameters

-vserver <vserver name> - Vserver (privilege: advanced)

Vserver name.

-domain <CIFS domain> - Domain (privilege: advanced)

The domain name of accountname.

-accountname <CIFS account> - User (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
```

vserver cifs superuser show

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you specify this parameter, the command displays superuser information of only the specified Vservers.

[-domain <CIFS domain>] - Domain (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 (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 (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
```

vserver cifs symlink create

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.

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.

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.

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
```

vserver cifs symlink delete

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/
```

vserver cifs symlink modify

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.

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.

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.

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
```

vserver cifs symlink show

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 Data 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.
```

vserver cifs users-and-groups update-names

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 (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 (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 (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.
```

vserver cifs users-and-groups local-group add-members

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
```

vserver cifs users-and-groups local-group create

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
```

vserver cifs users-and-groups local-group delete

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
```

vserver cifs users-and-groups local-group modify

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"
```

vserver cifs users-and-groups local-group remove-members

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
```

vserver cifs users-and-groups local-group rename

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
```

vserver cifs users-and-groups local-group show-members

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.
```

vserver cifs users-and-groups local-group show

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.
```

vserver cifs users-and-groups local-user create

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:
```

vserver cifs users-and-groups local-user delete

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
```

vserver cifs users-and-groups local-user modify

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
```

vserver cifs users-and-groups local-user rename

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
```

vserver cifs users-and-groups local-user set-password

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.
```

vserver cifs users-and-groups local-user show-membership

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.
```

vserver cifs users-and-groups local-user show

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.
```

vserver cifs users-and-groups privilege add-privilege

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
```

vserver cifs users-and-groups privilege remove-privilege

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
```

vserver cifs users-and-groups privilege reset-privilege

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.
```

vserver cifs users-and-groups privilege show

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
```

vserver config-replication commands

vserver config-replication pause

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 configuration baseline. This command aborts the ongoing baseline generation activity, unlocks the Vserver and temporarily pauses configuration replication for the Vserver. Command confirmations has 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 is 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 `vserver config replication resume` command.

Parameters

-vserver <vserver name> - Vserver name (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
```

vserver config-replication resume

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 `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 (privilege: advanced)

== Examples

```
cluster::> vserver config replication resume -vserver vs1
```

vserver config-replication show

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 ]` }

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

[-vserver <vserver>] - Vserver (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 (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
```

vserver data-policy commands

vserver data-policy export

Display a data policy

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

Description

The vserver-data policy export command displays the current data policy for a Vserver with Infinite Volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver with Infinite Volume for which the data policy will be displayed.

Examples

The following example shows the current data policy.

```
cluster1::> vserver data-policy export -vserver vs1

{ "ruleset_format_version" : "1.0",
  "rules" : [
    { "rule_label" : "default",
      "rule_id" : "ec17a05f-7785-11e1-baf4-123478563412",
      "rule_scope" : [],
      "rule_epoch" : { "epoch_reference" : "ctime" },
      "rule_epochs" : {
        "0" : {
          "local" : {
            "metadata" : {
              "storageservice" : "-"
            }
          }
        }
      }
    }
  ]
}
```

vserver data-policy import

Import a data policy

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

Description

The vserver data-policy import command sets a new data policy for a Vserver with Infinite Volume. After entering the command, you are prompted to type or paste the content of the new data policy. When you are done, press ENTER on a blank line.

Parameters

`-vserver <vserver name>` - Vserver Name

This specifies the Vserver with Infinite Volume for which the data policy will be changed.

Examples

The following examples attempt to change the Vserver data policy, first with bad content, and then with an acceptable data policy.

```
cluster1::> vserver data-policy import -vserver vs1

Enter the contents of the file data policy for Vserver "vs1":
Press <Enter> when done

{ "foo" : "bar" }
Error: command failed: Data Policy validation failed:
'ruleset_format_version'
       is a required field.

cluster1::> vserver data-policy import -vserver vs1

Enter the contents of the file data policy for Vserver "vs1":
Press <Enter> when done

{ "ruleset_format_version" : "1.0",
  "rules" : [
    { "rule_label" : "default",
      "rule_id" : "ec17a05f-7785-11e1-baf4-123478563412",
      "rule_scope" : [],
      "rule_epoch" : { "epoch_reference" : "ctime" },
      "rule_epochs" : {
        "0" : {
          "local" : {
            "metadata" : {
              "storageservice" : "-"
            }
          }
        }
      }
    }
  ]
}
```

vserver data-policy validate

Validate a data policy without import

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

Description

The vserver data-policy validate command checks a data policy for errors, without modifying the data policy for the Vserver with Infinite Volume.

Parameters

-vserver <vserver name> - Vserver Name

This specifies the Vserver with Infinite Volume for which the data policy will be validated.

Examples

The following examples show first a problem with a given data policy, and then an example of a valid data policy.

```
cluster1::> vserver data-policy validate -vserver vs1

Enter the contents of the file data policy for Vserver "vs1":
Press <Enter> when done

{ "foo" : "bar" }
Error: command failed: Data Policy validation failed:
'ruleset_format_version'
       is a required field.

cluster1::> vserver data-policy validate -vserver vs1

Enter the contents of the file data policy for Vserver "vs1":
Press <Enter> when done

{ "ruleset_format_version" : "1.0",
  "rules" : [
    { "rule_label" : "default",
      "rule_id" : "ec17a05f-7785-11e1-baf4-123478563412",
      "rule_scope" : [],
      "rule_epoch" : { "epoch_reference" : "ctime" },
      "rule_epochs" : {
        "0" : {
          "local" : {
            "metadata" : {
              "storageservice" : "-"
            }
          }
        }
      }
    }
  ]
}

Data Policy validation succeeded: No errors found.
```

vserver export-policy commands

vserver export-policy check-access

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. This command is not supported on Infinite 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 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
- *flexcache* - The FlexCache protocol

-access-type {read|read-write} - 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}] - 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.

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
Path                          Policy    Owner     Owner Type  Index
Access
----------------------------- ---------- --------- ---------- ------
----------
/                             default    vs1_root  volume          1 read
/dir1                         default    vs1_root  volume          1 read
/dir1/dir2                    default    vs1_root  volume          1 read
/dir1/dir2/flex1              data       flex_vol  volume         10 read
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
Path                          Policy    Owner     Owner Type  Index
Access
----------------------------- ---------- --------- ---------- ------
----------
/                             default    vs1_root  volume          1 read
/dir1                         default    vs1_root  volume          1 read
/dir1/dir2                    default    vs1_root  volume          1 read
/dir1/dir2/flex1              data       flex_vol  volume         10 read-
```

```
write
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
Path                          Policy    Owner     Owner Type  Index
Access
----------------------------- --------- --------- ---------- ------
----------
/                             default   vs1_root  volume          1 read
/dir1                         default   vs1_root  volume          1 read
/dir1/dir2                    default   vs1_root  volume          1 read
/dir1/dir2/flex1              data      flex_vol  volume         10 read
/dir1/dir2/flex1/qt1          primarynames
                                        qt1       qtree           0
denied
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
Path                          Policy    Owner     Owner Type  Index
Access
----------------------------- --------- --------- ---------- ------
----------
/                             default   vs1_root  volume          1 read
/dir1                         default   vs1_root  volume          1 read
/dir1/dir2                    default   vs1_root  volume          1 read
/dir1/dir2/flex1              data      flex_vol  volume         10 read
/dir1/dir2/flex1/qt1          primarynames
                                        qt1       qtree           2
denied
5 entries were displayed.
```

vserver export-policy copy

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
```

vserver export-policy create

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 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

- vserver export-policy rule create

vserver export-policy delete

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
```

vserver export-policy rename

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
```

vserver export-policy show

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.
```

vserver export-policy access-cache config modify-all-vservers

Modify exports access cache configuration for all Vservers

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

This command is not supported in a cluster with effective cluster version of Data ONTAP 9.0.0 or later. The access cache settings are modified on a per-Vserver basis starting Data ONTAP 9.0.0. See the vserver export-policy access-cache config modify command.

Parameters

`[-ttl-positive <integer>]` - TTL For Positive Entries (Secs) (privilege: advanced)

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) (privilege: advanced)

This parameter specifies the duration after which negative access cache entries will be refreshed when the client accesses.

`[-harvest-timeout <integer>]` - Harvest Timeout (Secs) (privilege: advanced)

This parameter specifies the time period after which Data ONTAP deletes unused entries in the access cache.

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 Data ONTAP 9.0.0.

```
cluster1::*> vserver export-policy access-cache config modify-all-vservers
-ttl-positive 36000 -ttl-negative 3600 -harvest-timeout 43200

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

- vserver export-policy access-cache config modify

vserver export-policy access-cache config modify

Modify exports access cache configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* 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 (privilege: advanced)

This parameter specifies the Vserver name for which the timeout values need to be modified.

`[-ttl-positive <integer>]` - TTL For Positive Entries (Secs) (privilege: advanced)

This parameter specifies the duration after which positive access cache entries will be refreshed upon client access.

`[-ttl-negative <integer>]` - TTL For Negative Entries (Secs) (privilege: advanced)

This parameter specifies the duration after which negative access cache entries will be refreshed upon client access.

`[-harvest-timeout <integer>]` - Harvest Timeout (Secs) (privilege: advanced)

This parameter specifies the time period after which Data ONTAP deletes unused entries in the access cache.

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

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
```

vserver export-policy access-cache config show-all-vservers

Display exports access cache configuration for all Vservers

Availability: This command is available to *cluster* administrators at the *advanced* 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. Data 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 Data ONTAP does not use an entry that is stored in the access cache for this period of time, it deletes the entry.

This command is not supported in a cluster with effective cluster version of Data ONTAP 9.0.0 or later. The access cache settings are stored on a per-Vserver basis starting Data ONTAP 9.0.0. See the 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 Data 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

- vserver export-policy access-cache config show

vserver export-policy access-cache config show

Display exports access cache configuration

Availability: This command is available to *cluster* and *Vserver* administrators at the *advanced* 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 Data 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 (privilege: advanced)

If this parameter is specified, the command displays the timeout values for the specified Vserver.

[-ttl-positive <integer>] - TTL For Positive Entries (Secs) (privilege: advanced)

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) (privilege: advanced)

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) (privilege: advanced)

If this parameter is specified, the command displays the timeout values for Vservers whose harvest-timeout matches the provided 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
               (secs)        (secs)      (secs)          (secs)
-------- ------------- ------------ ----------- ---------------
vs0               300           60           1            3600
vs1             36000         3600           5            3600
2 entries were displayed.
```

vserver export-policy cache flush

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
```

vserver export-policy config-checker show

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

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|Reschedule|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
```

vserver export-policy config-checker start

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.

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.
```

vserver export-policy config-checker stop

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.

> 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
```

vserver export-policy config-checker rule delete

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 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

- vserver export-policy config-checker rule show

vserver export-policy config-checker rule show

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
```

vserver export-policy netgroup check-membership

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. Data 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 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

- vserver export-policy netgroup queue show

vserver export-policy netgroup cache show

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
```

vserver export-policy netgroup queue show

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         -
```

vserver export-policy rule add-clientmatches

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 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. 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

- vserver export-policy rule show

vserver export-policy rule create

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 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.

`[-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
- *flexcache* - The FlexCache 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 FlexGroups 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. 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.

  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.

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.

  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.

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. NFS clients typically associate user ID 65534 with the user name nobody. In clustered Data ONTAP, this user ID is associated with the user pcuser. 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.

  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.

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 (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 (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 ownership of files that they own.

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

- vserver export-policy create

vserver export-policy rule delete

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 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 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

- vserver export-policy rule setindex
- vserver export-policy rule show

vserver export-policy rule modify

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 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 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
- *flexcache* - The FlexCache 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 FlexGroups 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. 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.

  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.

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.

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.

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. 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.

[-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.

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.

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 (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 (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 can change ownership of files that they own.

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

- vserver export-policy rule setindex
- vserver export-policy rule show

vserver export-policy rule remove-clientmatches

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 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

- vserver export-policy rule show

vserver export-policy rule setindex

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 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

- vserver export-policy rule show

vserver export-policy rule show

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
- *flexcache* - The FlexCache 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.

  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.

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.

  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.

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.

  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.

  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.

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 (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 (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 (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 ownership of the files that they own. If the Vserver NTFS 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 (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.
```

vserver fcp commands

vserver fcp create

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 (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.

Unles 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 (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* .

vserver fcp delete

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 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

- vserver fcp modify

vserver fcp modify

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 (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.

Unles 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 (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`* .

vserver fcp show

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.

vserver fcp start

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 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

- system license add

vserver fcp stop

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* .

vserver fcp initiator show

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

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 <text>] - Initiator WWPN

Use this parameter to display the FCP initiators that matches the World Wide Port Name (WWPN) that you specify.

[-port-address <Hex Integer>] - Port Address

Use this parameter to display FCP initiators that match the port address that you specify.

[-wwnn <text>] - Initiator WWNN

Use this parameter to display the FCP initiator that matches the World Wide Node Name (WWNN) 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.

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.

vserver fcp interface show

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.

vserver fcp nameserver show

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 <integer>] - Port Identifier

Use this parameter to select the assigned port identifier of the LIF.

[-unzoned <true>] - Show unzoned name server entires

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.
```

vserver fcp ping-igroup show

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 (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.
```

vserver fcp ping-initiator show

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 (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.
```

vserver fcp portname set

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 network interface modify to change the administration status of the logical interface.

Parameters

`-vserver <Vserver Name>` - Vserver (privilege: advanced)
Specifies the Vserver.

`-lif <lif-name>` - Logical Interface (privilege: advanced)
Specifies the logical interface to which you want to assign a new WWPN.

`-wwpn <text>` - FCP Adapter WWPN (privilege: advanced)
Specifies the WWPN that you want to change.

`[-f, -force <true>]` - Force (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

- network interface modify

vserver fcp portname show

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.

vserver fcp topology show

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.

vserver fcp wwn blacklist show

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] }

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

[-wwn <text>] - World Wide Name (privilege: advanced)

Selects the blacklisted WWNs that match the parameter value.

[-vserver <Vserver Name>] - Vserver Name (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.

vserver fcp wwpn-alias remove

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* .

vserver fcp wwpn-alias set

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* .

vserver fcp wwpn-alias show

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).

You can also use these commands to display WWPN aliases:

- lun igroup show
- lun igroup create
- lun igroup add
- lun igroup remove
- 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

- lun igroup show
- lun igroup create
- lun igroup add
- lun igroup remove
- vserver fcp show

vserver fpolicy commands

vserver fpolicy disable

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver fpolicy enable

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver fpolicy engine-connect

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

The FPolicy event and external engine must be attached to the FPolicy policy.

The FPolicy policy should be enabled.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy engine-disconnect

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy prepare-to-downgrade

Restore the FPolicy configuration to Earlier Release of Data ONTAP

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

Description

The `vserver fpolicy prepare-to-downgrade` command restores the FPolicy configurations for Data ONTAP based on the input parameter disable-feature-set.

Parameters

`-disable-feature-set <downgrade version>` - Data ONTAP Version (privilege: advanced)

This parameter specifies the Data ONTAP version that introduced the new FPolicy features and needs to be restored. The value can be one of the following:

- 9.0.0 - Disables the FPolicy features introduced in Data ONTAP release 9.0.0. These features include:
- FPolicy filters for setattr operation.
- FPolicy filter `exclude-directory` for directory operations.
- FPolicy Async Resiliency feature.
- FPolicy "Monitor Objects With No Extension" feature.

Examples

```
cluster1::*> vserver fpolicy prepare-to-downgrade -disable-feature-set
9.0.0
```

vserver fpolicy show-enabled

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy show-engine

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.

This command is not supported for a Vserver with Infinite Volume.

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:
```

vserver fpolicy show-passthrough-read-connection

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
```

vserver fpolicy show

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver fpolicy policy create

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.

This command is not supported for a Vserver with Infinite Volume.

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 Data 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* .

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
```

vserver fpolicy policy delete

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy policy modify

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.

This command is not supported for a Vserver with Infinite Volume.

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 Data 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* .

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
```

vserver fpolicy policy show

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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.
```

vserver fpolicy policy event create

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.

This command is not supported for a Vserver with Infinite Volume.

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.

> 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.

> 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.

If you specify a value for the `-filters` parameter, then you must also specify valid values for the `-file-operations` and `-protocol` parameters.

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*.

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
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
             Volume Operation: true
```

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
```

vserver fpolicy policy event delete

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.

> This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy policy event modify

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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.

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.

If you specify a value for the `-filters` parameter, then you must also specify valid values for the `-file-operations` and `-protocol` parameters.

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*.

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
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
            Volume Operation: true
```

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
```

vserver fpolicy policy event show

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

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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
```

vserver fpolicy policy external-engine create

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.

This command is not supported for a Vserver with Infinite Volume.

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 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 security certificate install command with `-type` set to *`server`* .

`[-reqs-cancel-timeout <[<integer>h][<integer>m][<integer>s]>]` - Timeout for Canceling a Request (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>h][<integer>m][<integer>s]>]` - Timeout for Aborting a Request (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>h][<integer>m][<integer>s]>]` - Interval for Sending Status Requests (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 (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 (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. By default, it is 50.

`[-server-progress-timeout <[<integer>h][<integer>m][<integer>s]>]` - Timeout for Disconnecting Non-responsive Server (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>h][<integer>m][<integer>s]>] - Interval for Sending Keep-Alive Messages (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 (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 (privilege: advanced)

This parameter specifies the send 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 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-defualt value to set the size (in bytes) of the send buffer.

[-session-timeout <[<integer>h][<integer>m][<integer>s]>] - Session ID Purge Timeout During Reconnection (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>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.

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 -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
SSL Option for External Communication: mutual-auth
           FQDN or Custom Common Name: Sample1-FPolicy-Client
                        Serial Number: 8DDE112A114D1FBC
                  Certificate Authority: TASample1
```

Related Links

- security certificate install

vserver fpolicy policy external-engine delete

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy policy external-engine modify

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.

This command is not supported for a Vserver with Infinite Volume.

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 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 security certificate install command with `-type` set to *`server`*.

`[-reqs-cancel-timeout <[<integer>h][<integer>m][<integer>s]>]` - Timeout for Canceling a Request (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>h][<integer>m][<integer>s]>]` - Timeout for Aborting a Request (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>h][<integer>m][<integer>s]>]` - Interval for Sending Status Requests (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 (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 (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. By default, it is 50.

`[-server-progress-timeout <[<integer>h][<integer>m][<integer>s]>]` - Timeout for Disconnecting Non-responsive Server (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>h][<integer>m][<integer>s]>]` - Interval for Sending Keep-Alive Messages (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 (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 (privilege: advanced)

This parameter specifies the send 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 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-defualt value to set the size (in bytes) of the send buffer.

[-session-timeout <[<integer>h][<integer>m][<integer>s]>] - Session ID Purge Timeout During Reconnection (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>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.

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
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

- security certificate install

vserver fpolicy policy external-engine show

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

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.

This command is not supported for a Vserver with Infinite Volume.

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>h][<integer>m][<integer>s]>] - Timeout for Canceling a Request (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>h][<integer>m][<integer>s]>] - Timeout for Aborting a Request (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>h][<integer>m][<integer>s]>] - Interval for Sending Status Requests (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 (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 (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>h][<integer>m][<integer>s]>] - Timeout for Disconnecting Non-responsive Server (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>h][<integer>m][<integer>s]>] - Interval for Sending Keep-Alive Messages (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 (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 (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>h][<integer>m][<integer>s]>] - Session ID Purge Timeout During Reconnection (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>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.

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
 Vserver         Engine      Servers           Servers             Port
Engine Type
 --------------- ----------  ----------------- ----------------- ------
----------
 Cluster         cserver_eng  9.9.9.9          -                     9
synchronous
 vs1.example.com cserver_eng  9.9.9.9          -                     9
synchronous
 vs1.example.com v1n1         1.1.1.1          2.2.2.2               1
synchronous
 vs2.example.com cserver_eng  9.9.9.9          -                     9
synchronous
 vs2.example.com v2n1         3.3.3.3          5.5.5.5               2
synchronous
 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
           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
```

vserver fpolicy policy scope create

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.

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.

This command is not supported for a Vserver with Infinite Volume.

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 "*".

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 "*".

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 "*".

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 "?".

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 (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 *`false`* .

[-is-monitoring-of-objects-with-no-extension-enabled {true|false}] - Is Monitoring of Objects with No Extension Enabled (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* .

> 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.
```

vserver fpolicy policy scope delete

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.

> This command is not supported for a Vserver with Infinite Volume.

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
```

vserver fpolicy policy scope modify

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.

This command is not supported for a Vserver with Infinite Volume.

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 "*".

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 "*".

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 "*".

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 "?".

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 (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 *`false`* .

`[-is-monitoring-of-objects-with-no-extension-enabled {true|false}]` - Is Monitoring of Objects with No Extension Enabled (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`* .

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.
```

vserver fpolicy policy scope show

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 extention 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.

This command is not supported for a Vserver with Infinite Volume.

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 (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 (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.
```

vserver group-mapping commands

vserver group-mapping create

Create a group mapping

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

Description

The `vserver group-mapping create` command creates a group mapping. Group mappings are applied in the order in which they occur in the priority list; for example, a group mapping that occurs at position 2 in the priority list is applied before a group mapping that occurs at position 3. Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX, and UNIX-to-Windows) has its own priority list. Data ONTAP prevents you from creating two group 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 1024 group mappings in each direction.

The `vserver group-mapping create` command is not supported on Vservers with FlexVol volumes.

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 group mapping.

-direction {krb-unix|win-unix|unix-win} - Name Mapping Direction

This parameter specifies the direction of the group mapping. Possible values are *krb-unix* for a Kerberos-to-UNIX group mapping, *win-unix* for a Windows-to-UNIX group mapping, and *unix-win* for a UNIX-to-Windows group mapping.

-position <integer> - Position

This parameter specifies the group mapping's position in the priority list. Specify the position as a positive integer.

If you want to create a new group mapping at a position that is already occupied in the priority list, use the vserver group-mapping insert command instead of the `vserver group-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.

Examples

The following example creates a group 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 Groups.

```
cluster1::> vserver group-mapping create -vserver vs1 -direction unix-win
-position 5 -pattern cifs -replacement "EXAMPLE\\Domain Groups"
```

Related Links

- vserver group-mapping insert

vserver group-mapping delete

Delete a group mapping

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

Description

The `vserver group-mapping delete` command deletes a group mapping.

The `vserver group-mapping delete` command is not supported on Vservers with FlexVol volumes.

Parameters

`-vserver <vserver name>` - Vserver

This parameter specifies the Vserver from which you want to delete the group mapping.

`-direction {krb-unix|win-unix|unix-win}` - Name Mapping Direction

This parameter specifies the direction of the group mapping that you want to delete.

`-position <integer>` - Position

This parameter specifies the position of the group mapping that you want to delete. Specify the position as a positive integer.

Examples

The following example deletes a group mapping on a Vserver named vs1. The group mapping is from UNIX to Windows and is at position 5.

```
cluster1::> vserver group-mapping delete -vserver vs1 -direction unix-win
-position 5
```

vserver group-mapping insert

Create a group mapping at a specified position

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

Description

The `vserver group-mapping insert` command creates a group 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, and UNIX-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 1024 group mappings in each direction.

The `vserver group-mapping insert` command is not supported on Vservers with FlexVol volumes.

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 group mapping.

-direction {krb-unix|win-unix|unix-win} - Name Mapping Direction

This parameter specifies the direction of the group mapping. Possible values are *krb-unix* for a Kerberos-to-UNIX group mapping, *win-unix* for a Windows-to-UNIX group mapping, and *unix-win* for a UNIX-to-Windows group mapping.

-position <integer> - Position

This parameter specifies the position in the priority list at which you want to insert the new group 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.

Examples

The following example creates a group mapping on a Vserver named vs1. It is a group mapping from Kerberos to UNIX. It is inserted into the priority list at position 2. The group mapping maps any principal in the Kerberos realm SEC.EXAMPLE.COM to the UNIX group name corresponding to the principal's base name with any instance names removed; for example, artists/admin@SEC.EXAMPLE.COM is mapped to artists.

```
cluster1::> vserver group-mapping insert -vserver vs1 -direction krb-unix
-position 2 -pattern "([^@/]+)(/[^@]+)?@SEC.EXAMPLE.COM" -replacement "\1"
```

vserver group-mapping modify

Modify a group 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 group-mapping modify` command modifies the pattern, the replacement pattern, or both of a specified group 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 1024 group mappings in each direction.

The `vserver group-mapping modify` command is not supported on Vservers with FlexVol volumes.

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 group mapping.

-direction {krb-unix|win-unix|unix-win} - Name Mapping Direction

This parameter specifies the direction of the group mapping. Possible values are *krb-unix* for a Kerberos-to-UNIX group mapping, *win-unix* for a Windows-to-UNIX group mapping, and *unix-win* for a UNIX-to-Windows group mapping.

-position <integer> - Position

This parameter specifies the group mapping's position in the priority list. A position is specified as a positive integer. Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX, and UNIX-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.

Examples

The following example modifies the group mapping on the Vserver named vs1 and direction win-unix, at position 3. The pattern to be matched is changed to "EXAMPLE\(.+)".

```
cluster1::> vserver group-mapping modify -vserver vs1 -direction win-unix
-position 3 -pattern "EXAMPLE\\(.+)"
```

vserver group-mapping show

Display group mappings

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

Description

The `vserver group-mapping show` command displays information about group 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 group 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 group mappings, run the command with the `-direction krb-unix` parameter.

The `vserver group-mapping show` command is not supported on Vservers with FlexVol 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 you specify this parameter, the command displays information only about the group mapping or mappings that match the specified Vserver.

[-direction {krb-unix|win-unix|unix-win}] - Name Mapping Direction

If you specify this parameter, the command displays information only about the group mapping or mappings that have the specified mapping direction.

[-position <integer>] - Position

If you specify this parameter, the command displays information only about the group 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 group 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 group mapping or mappings that use the specified replacement pattern.

Examples

The following example displays information about all group mappings:

```
cluster1::> vserver group-mapping show
Vserver        Direction Position
-------------- --------- --------
vs1            win-unix  1        Pattern: EXAMPLE\\artists
                               Replacement: nobody
vs1            unix-win  1        Pattern: EXAMPLE\\(.+)
                               Replacement: \_1
vs2            win-unix  1        Pattern: (.+)
                               Replacement: EXAMPLE\\artists
```

vserver group-mapping swap

Exchange the positions of two group mappings

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

Description

The `vserver group-mapping swap` command exchanges the positions of two group mappings in the priority list.

The `vserver group-mapping swap` command is not supported on Vservers with FlexVol volumes.

Parameters

-vserver <vserver name> - Vserver

This parameter specifies the Vserver on which the group mappings are located.

-direction {krb-unix|win-unix|unix-win} - Name Mapping Direction

This parameter specifies the direction of the group 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 group mapping that you want to exchange. Specify a position as a positive integer.

-with-position <integer> - Position of an existing group mapping entry in the list of group 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 group mapping that you want to exchange. Specify a position as a positive integer.

Examples

The following example exchanges the positions of two group mappings on a Vserver named vs1. The group mappings have the direction Windows-to-UNIX. The group mappings are exchanged between positions 2 and 4.

```
cluster1::> vserver group-mapping swap -vserver vs1 -direction win-unix
-position 2 -with-position 4
```

vserver iscsi commands

vserver iscsi create

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 (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.

[-max-error-recovery-level <integer>] - Max Error Recovery Level (privilege: advanced)

Specifies the maximum error recovery level allowed by the iSCSI service. You can specify 0, 1, or 2, or you can accept the default. The default is zero. The actual error recovery level depends on the negotiated error recovery level between the initiator and the iSCSI target when the session is created.

- 0 - Session failure recovery
- 1 - Digest failure recovery
- 2 - Connection failure recovery

[-retain-timeout <integer>] - RFC3720 DefaultTime2Retain Value (in sec) (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) (privilege: advanced)

Specifies the login phase duration. The default is 15 seconds.

[-max-conn-per-session <integer>] - Max Connections per Session (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 (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) (privilege: advanced)

Specifies the TCP receive window size (in bytes). The default is 131,400 bytes.

[-f, -force <true>] - Allow Non-Vendor Target Name (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.

vserver iscsi delete

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.

> You must first disable the service with the command 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

- vserver iscsi modify

vserver iscsi modify

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 (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 "_".

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.

[-max-error-recovery-level <integer>] - Max Error Recovery Level (privilege: advanced)

Specifies the maximum error recovery level the iSCSI service negotiates with iSCSI initiators during login phase.

- 0 - Session failure recovery
- 1 - Digest failure recovery
- 2 - Connection failure recovery

[-retain-timeout <integer>] - RFC3720 DefaultTime2Retain Value (in sec) (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) (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 (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 (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) (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 (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.

vserver iscsi show

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.

[-max-error-recovery-level <integer>] - Max Error Recovery Level (privilege: advanced)

Selects the iSCSI services with a maximum error recovery level that matches the parameter value.

[-retain-timeout <integer>] - RFC3720 DefaultTime2Retain Value (in sec) (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) (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 (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 (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) (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
                Max Error Recovery Level: 0
      DefaultTime2Retain Value (in sec): 20
          Login Phase Duration (in sec): 20
            Max Connections per Session: 4
                     Max I/O per Session: 128
TCP Window Size all Sessions (in bytes): 131400
1 entries were displayed.
```

Displays the detailed entries for all entries.

vserver iscsi start

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 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

- vserver iscsi modify

vserver iscsi stop

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 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

- vserver iscsi modify

vserver iscsi command show

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

- "lmm_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.

vserver iscsi connection show

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.

vserver iscsi connection shutdown

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 vserver iscsi session shutdown command.

Parameters

-vserver <Vserver Name> - Vserver (privilege: advanced)

Specifies the Vserver.

-tpgroup <text> - Target Portal Group (privilege: advanced)

Specifies the target portal group that contains the connection you want to shut down.

-tsih <integer> - Target Session ID (privilege: advanced)

Specifies the target session ID that you want to shut down.

-connection-id <integer> - Connection ID (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

- vserver iscsi session shutdown

vserver iscsi initiator show

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.

vserver iscsi interface disable

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.

vserver iscsi interface enable

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.

vserver iscsi interface modify

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 (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
```

vserver iscsi interface show

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 (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.

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
```

vserver iscsi interface accesslist add

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 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 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.

ⓘ 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

- vserver iscsi interface disable
- vserver iscsi interface accesslist remove

vserver iscsi interface accesslist remove

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.

vserver iscsi interface accesslist show

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.

vserver iscsi isns create

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.

> 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 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.

> 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

- network interface create
- network routing-groups route create

- network routing-groups route show

vserver iscsi isns delete

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.

vserver iscsi isns modify

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.

> 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.

Related Links

- network routing-groups route create
- network routing-groups route show

vserver iscsi isns show

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.

vserver iscsi isns start

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.

vserver iscsi isns stop

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.

vserver iscsi isns update

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.

vserver iscsi security add-initator-address-ranges

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.

vserver iscsi security create

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 -initiator
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.

vserver iscsi security default

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 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 -security 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

- vserver iscsi security create

vserver iscsi security delete

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
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.

vserver iscsi security modify

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 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
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

- vserver iscsi security delete

vserver iscsi security prepare-to-downgrade

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 Data 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 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 Data ONTAP.

Related Links

- vserver iscsi security show

vserver iscsi security remove-initator-address-ranges

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.

vserver iscsi security show

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. Data 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.

vserver iscsi session show

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 vserver iscsi connection show command to display connection information. Use the 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 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

- vserver iscsi connection show
- vserver iscsi session parameter show
- vserver iscsi modify

vserver iscsi session shutdown

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 vserver iscsi connection shutdown command.

Parameters

-vserver <Vserver Name> - Vserver (privilege: advanced)

Specifies the Vserver.

-tpgroup <text> - Target Portal Group (privilege: advanced)

Specifies the target portal group that contains the session you want to shutdown.

-tsih <integer> - Target Session ID (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

- vserver iscsi connection shutdown

vserver iscsi session parameter show

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 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

- vserver iscsi modify

vserver locks commands

vserver locks break

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 (privilege: advanced)
This parameter specifies the Vserver containing the lock.

`-volume <volume name>` - Volume (privilege: advanced)
This parameter specifies the name of the volume containing the lock.

`-lif <lif-name>` - Logical Interface (privilege: advanced)
This parameter specifies the logical interface through which the lock was established.

`-path <text>` - Object Path (privilege: advanced)
This parameter specifies a path to the lock.

| `-lockid <UUID>` - Lock UUID (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 (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 (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 (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 (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.

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.
```

vserver locks show

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

Parameters

{ [-fields <fieldname>,…]

If you specify 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.

[-lif <lif-name>] - Logical Interface

If you specify this parameter, the command displays information about locks established through the specified logical interface.

[-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. This field is *`false`* for constituents of an Infinite Volume.

[-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. Please 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.

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.
```

vserver migrate commands

vserver migrate cleanup

Remove migrating entity

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

Description

This command cleans up the migrate destination Vserver. If the Vserver migrate operation has failed or is paused, use this command to delete the destination Vserver.

Parameters

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Name of the Vserver which is being migrated.

Examples

cluster1::> vserver migrate cleanup -vserver test

vserver migrate cutover

Perform Cutover of the migrate operation

Availability: This command is available to *cluster* administrators at the *advanced* 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 (privilege: advanced)

Name of the Vserver which is being migrated.

Examples

cluster1::> vserver migrate cutover -vserver test

vserver migrate pause

Pause a Vserver migrate operation

Availability: This command is available to *cluster* administrators at the *advanced* 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 (privilege: advanced)

Name of the Vserer whose migrate will be paused.

Examples

cluster1::> vserver migrate pause -vserver test

vserver migrate repeer

Repeer Existing Vserver Peer Relationships after Vserver Migration

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

Description

This command re-establishes the existing Vserver peer relationships after Vserver migration has finished.

Parameters

`-source-cluster <text>` - Source Cluster Name (privilege: advanced)

Name of the source cluster.

`-vserver-name <text>` - migrated vserver name (privilege: advanced)

Name of the Vserver that has finished migrating.

Examples

cluster1::> vserver migrate repeer -source-cluster cluster-2 -vserver-name test.

vserver migrate resume

Resume a migrate operation

Availability: This command is available to *cluster* administrators at the *advanced* 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 (privilege: advanced)

Name of the Vserver being migrated.

`[-force {true|false}]` - Force flag for continuing with disruptive migrate (privilege: advanced)

If set to true, the Vserver migrate will be disruptive, and will continue to completion even if cutover lasts longer than its normal 30 second window.

[-auto-cutover {true|false}] - Automatically cutover when ready (privilege: advanced)

This parameter is to specify if the Vserver migrate operation should cutover automatically when ready.

[-skip-performance-check {true|false}] - Skip checking iops requirement of volume on destination aggregates (privilege: advanced)

If set to true, the destination aggregates will not be checked to see if they meet the IOPS requirements.

Examples

cluster1::> vserver migrate resume -vserver test -force false -auto-cutover false -skip-performance-check false

vserver migrate show-progress

Display status of volumes in a migrating Vservers

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

Description

This command displays data transfer progress of all the volumes in migrating Vservers.

Parameters

{ [-fields <fieldname>,…]

This specifies the fields that need to be displayed.

| [-instance] }

If this parameter is specified, the command displays detailed volume progress information.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Name of the Vserver which is migrating.

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

If this parameter is specified, the command displays the details of volume progress that matches specified volume.

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

If this parameter is specified, the command displays the details of volume progress that matches specified Vserver UUID.

[-source-cluster-uuid <UUID>] - Source Cluster Uuid (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified source cluster UUID.

[-bytes-transferred {<integer>[KB|MB|GB|TB|PB]}] - Bytes transferred per volume (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified bytes transferred.

[-bytes-to-be-transferred {<integer>[KB|MB|GB|TB|PB]}] - Bytes to be transferred per volume (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified bytes to be transferred.

[-transfer-rate <text>] - Rate of data transfers (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified transfer rate.

[-transfer-start-time <integer>] - Transfer start time (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified transfer start time.

[-total-bytes-to-be-transferred {<integer>[KB|MB|GB|TB|PB]}] - Total bytes to be transferred per vserver (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified total bytes to be transferred.

[-progress-time-last-updated <integer>] - Real time progress time given by SnapMirror (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified real time progress time given by snapmirror.

[-last-transfer-time <integer>] - Last trasnfer time (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified last transfer time.

[-progress-bytes-last-updated {<integer>[KB|MB|GB|TB|PB]}] - Real time progress bytes given by SnapMirror (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified real time progress bytes given by snapmirror.

[-time-remaining <text>] - Remaining time for bytes transfer (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified time remaining for bytes transfer.

[-percent-complete <percent>] - Percentage of transfer completed (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified percentage of transfer completed.

[-total-bytes-transferred {<integer>[KB|MB|GB|TB|PB]}] - Total bytes trasferred per vserver (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified total bytes transferred per vserver.

[-average-transfer-rate <text>] - Average transfer rate per vserver (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified average transfer rate per vserver.

[-total-time-remaining <text>] - Total Time Remaining per vserver (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified total time remaining per vsrever.

[-total-percent-complete <percent>] - Total Percent complete per vserver (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified total percent of transfer complete per vserver.

[-total-used {<integer>[KB|MB|GB|TB|PB]}] - Total Used (privilege: advanced)

If this parameter is specified, the command displays the details of volume progress that matches specified total used size.

Examples

cluster1::> vserver migrate show-progress -vserver test

vserver migrate show-volume

Display status of volumes in a migrating Vservers

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

Description

This command displays status of all the volumes in a migrating Vservers.

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 (privilege: advanced)

Name of the Vserver which is migrating.

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

If this parameter is specified, the command displays detailed volume status information that matches specified volume.

[-volume-dsid <integer>] - Volume DSID (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified DSID.

[-volume-msid <integer>] - Volume MSID (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified MSID.

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

If this parameter is specified, the command displays the details of volume status that matches specified Vserver UUID.

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

If this parameter is specified, the command displays the details of volume status that matches specified node.

[-state {Transfer | ReadyForCutover | PreCutover | Penultimate | CutoverStarted | CutoverComplete | Cleanup | MigrateFailed}] - Status of the transfer (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified vserver migrate status.

[-last-transfer-duration <integer>] - Duration of last transfer in seconds (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified last transfer duration.

[-last-transfer-done {true|false}] - Is the last transfer done (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified last transfer done flag status.

[-break-done {true|false}] - Is the break done (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified break-done flag status.

[-errors <text>] - Errors in volume operation if any (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified errors.

[-last-transfer-queue-time <MM/DD/YYYY HH:MM:SS>] - Time of the last transfer (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified last transfer queue time.

[-transfer-completed {true|false}] - Is the transfer completed (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified transfer completed flag status.

[-cutover-transfer-count <integer>] - Number of transfers within cutover threshold (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified cutover transfer count.

[-force {true|false}] - Volume force cutover flag (privilege: advanced)

If this parameter is specified, the command displays the details of volume status that matches specified volume force cutover flag.

Examples

cluster1::> vserver migrate show-volume -vserver test

vserver migrate show

Display status of migrating Vservers

Availability: This command is available to *cluster* administrators at the *advanced* 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] (privilege: advanced)

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 (privilege: advanced)

Name of the Vserver which is migrating.

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

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver UUID.

[-transaction-id <integer>] - Transaction Id (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified transaction-id. A transaction-id is a 64 bit integer which identifies each Vserver migrate operation uniquely

[-destination-cluster <Cluster name>] - Destination Cluster Name (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified destination cluster.

[-source-cluster <Cluster name>] - Source Cluster Name (privilege: advanced)

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-triggerred|cutover-started|cutover-complete|migrate-paused|migrate-complete|migrate-complete-with-warnings|migrate-failed|post-cutover-cleanup|cleanup-failed|manual-cleanup|destination-cleaned-up|migrate-

pausing|cleanup-pausing}] - Vserver migrate status (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate start time.

[-completion-time <MM/DD/YYYY HH:MM:SS>] - Migrate operation finish time (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate completion time.

[-last-pause-time <MM/DD/YYYY HH:MM:SS>] - Last migrate pause time (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate last resume time.

[-last-rollback-time <MM/DD/YYYY HH:MM:SS>] - Last migrate rollback time (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate last rollback time.

[-cutover-trigger-time <MM/DD/YYYY HH:MM:SS>] - Cutover trigger time (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate cutover trigger time.

[-cutover-start-time <MM/DD/YYYY HH:MM:SS>] - Cutover start time (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate cutover completion time.

[-rollback-count <integer>] - Rollback Count (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate status details.

[-cutover-window <integer>] - cutover duration(seconds) (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate cutover window.

[-ipspace <IPspace>] - Destination cluster IPspace Name for vserver (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified ipspace.

[-force {true|false}] - Force cutover until completion (privilege: advanced)

If force is set to "true", the command will only display Vserver migrate operation information about Vservers where the value of force is set to "true". If set to "false", the command will only display Vserver migrate operation information where force is set to "false".

[-aggr-list <aggregate name>,…] - Aggregate list for creating the volumes (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified Vserver migrate type.

[-is-past-point-of-no-return {true|false}] - Indicate point of no return for migrate (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified past point of no return flag.

[-current-migrate-operation {none|start|resume|pause|cleanup|cutover}] - Current Migrate operation (privilege: advanced)

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}] - Last Migrate operation (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified local vserver id.

[-group-id <integer>] - Group ID (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified group id.

[-partner-vserver-id <integer>] - Partner Vserver ID (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified partner vserver id.

[-partner-group-id <integer>] - Partner group ID (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified partner group id.

[-auto-cutover {true|false}] - Automatic cutover (privilege: advanced)

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 (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified skip performance flag.

[-cutover-ready-max-transfer-time-limit <integer>] - Transfer duration for marking ready for cutover (seconds) (privilege: advanced)

If this parameter is specified, the command displays Vserver migrate information that matches the specified time required to mark cutover ready.

Examples

cluster1::> vserver migrate show -vserver test

vserver migrate start

Start the Vserver migrate operation

Availability: This command is available to *cluster* administrators at the *advanced* 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 (privilege: advanced)

Name of the Vserver which needs to be migrated.

-source-cluster <Cluster name> - Source Cluster Name (privilege: advanced)

Name of the source cluster.

[-check-only {true|false}] - Check if migrate can be done (privilege: advanced)

Runs the prechecks and tells if the Vserver migrate operation can be started or not.

[-ipspace <IPspace>] - Destination cluster IPspace Name for vserver (privilege: advanced)

Name of the IPspace in the destination cluster.

[-aggr-list <aggregate name>,…] - Aggregate list (privilege: advanced)

Provide the list of aggregates where the volumes will be created in the destination cluster.

[-force {true|false}] - Force flag for continuing with disruptive migrate (privilege: advanced)

The force parameter is set to true when the user wants Vserver migrate operation to continue to completion even though the 30 sec cutover duration is not met. When this parameter is used, it indicates that the admin wants to perform a disruptive migrate operation.

[-auto-cutover {true|false}] - Automatically cutover when ready (privilege: advanced)

This parameter is to specify if the Vserver migrate operation should cutover automatically when ready.

[-skip-performance-check {true|false}] - Skip checking iops requirement of volume on destination aggregates (privilege: advanced)

This parameter is set to true when the user wants the Vserver migrate operation to skip checking IOPS requirement of volume on destination aggregates.

Examples

cluster1::> vserver migrate start -vserver test -source-cluster cluster-22 -ipspace ips1 -check-only true -aggr-list aggr1,aggr2 -force false -auto-cutover false -skip-performance-check false

vserver name-mapping commands

vserver name-mapping create

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, and UNIX-to-Windows) has its own priority list. Data 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 1024 name mappings in each direction.

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} - 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, and *unix-win* for a UNIX-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.

If you want to create a new name mapping at a position that is already occupied in the priority list, use the 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 cifs -replacement "EXAMPLE\\Domain Users -address
10.238.33.245/24"
cluster1::> vserver name-mapping create -vserver vs1 -direction unix-win
-position 5 -pattern cifs -replacement "EXAMPLE\\Domain Users -hostname
google.com"
```

Related Links

- vserver name-mapping insert

vserver name-mapping delete

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}` - 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
```

vserver name-mapping insert

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, and UNIX-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 1024 name mappings in each direction.

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} - 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, and *unix-win* for a UNIX-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
```

vserver name-mapping modify

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 1024 name mappings in each direction.

> 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} - 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, and *unix-win* for a UNIX-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, and UNIX-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"
```

vserver name-mapping refresh-hostname-ip

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 (privilege: advanced)

This parameter specifies the Vserver for which the hostname lookup needs to be done.

`[-direction {krb-unix|win-unix|unix-win}]` - Name Mapping Direction (privilege: advanced)

This optional parameter specifies the direction of the name-mapping entry for the hostname lookup.

`[-hostname <text>]` - Hostname (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
```

vserver name-mapping show

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}] - 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
```

vserver name-mapping swap

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}` - 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
```

vserver nfs commands

vserver nfs create

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 (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 (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 `disabled` . This parameter is not supported for Vservers with Infinite Volume.

`[-udp {enabled|disabled}]` - UDP Protocol

This optional parameter specifies whether to enable NFS access over UDP. The default setting is `enabled` . This parameter is not supported for Vservers with Infinite Volume.

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 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 (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 (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 (privilege: advanced)

This optional parameter specifies whether Data 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 (privilege: advanced)

This optional parameter specifies whether Data 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 (privilege: advanced)

This optional parameter specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to `ignore` , Data 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` , Data 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 (privilege: advanced)

This optional parameter specifies whether ownership of a file can be changed by superusers or by non-root users who currently own the file. If you set this parameter to `restricted` , the ownership of a file can be changed by superusers only. If you set this parameter to `unrestricted` , the ownership of a file can be changed by superusers and the current owner of the file. If you set this parameter to `use-export-policy` , the ownership of a file can be changed in accordance with the relevant export rules. The default setting is `use-export-policy` .

[-trace-enabled {true|false}] - NFS Response Trace Enabled (privilege: advanced)

This optional parameter specifies whether Data 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) (privilege: advanced)

This optional parameter specifies the amount of time, in seconds, after which Data 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) (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) (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.

Setting the parameter value greater than 65536 may cause performance degradation for existing connections using smaller values. Contact technical support for guidance.

[-v3-tcp-max-read-size <integer>] - NFSv3 TCP Maximum Read Size (bytes) (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 read requests. The range is 8192 to 1048576. The default setting is 65536 when created.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-v3-tcp-max-write-size <integer>] - NFSv3 TCP Maximum Write Size (bytes) (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 write requests. The range is 8192 to 65536. The default setting is 65536 when created.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-v4.0-acl {enabled|disabled}] - NFSv4.0 ACL Support

This optional parameter specifies whether Data ONTAP supports NFSv4.0 access control lists (ACLs). The default setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4.0-read-delegation {enabled|disabled}] - NFSv4.0 Read Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4.0 read delegations. The default setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4.0-write-delegation {enabled|disabled}] - NFSv4.0 Write Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4.0 write delegations. The default

setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4-fsid-change {enabled|disabled}] - Show Change in FSID as NFSv4 Clients Traverse Filesystems (privilege: advanced)

This optional parameter specifies whether Data ONTAP shows changes in file system identifiers (FSIDs) as NFSv4 clients traverse file systems. The default setting is `enabled` .

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 (privilege: advanced)

This optional parameter specifies whether Data 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. This parameter is not supported for Vservers with Infinite Volume.

[-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 normally taken from the NIS domain or the DNS domain in use. 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 (privilege: advanced)

This optional parameter specifies whether Data 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) (privilege: advanced)

This optional parameter specifies the time period in which Data 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 Data 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 clients. The default setting is `disabled` .

[-rquota {enabled|disabled}] - Rquota Enable

This optional parameter specifies whether to enable rquota over NFS. The default setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-implementation-domain <nfs domain>] - NFSv4.1 Implementation ID Domain (privilege: advanced)

This optional parameter specifies the NFSv4.1 implementation domain.

[-v4.1-implementation-name <text>] - NFSv4.1 Implementation ID Name (privilege: advanced)

This optional parameter specifies the NFSv4.1 implementation name.

[-v4.1-implementation-date <Date>] - NFSv4.1 Implementation ID Date (privilege: advanced)

This optional parameter specifies the NFSv4.1 implementation date.

[-v4.1-pnfs {enabled|disabled}] - NFSv4.1 Parallel NFS Support

This optional parameter specifies whether Data ONTAP supports parallel NFS over NFSv4.1. The default setting is `enabled` .

[-v4.1-referrals {enabled|disabled}] - NFSv4.1 Referral Support (privilege: advanced)

This optional parameter specifies whether Data ONTAP supports NFSv4.1 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 referrals, set this option to `disabled` ; otherwise, those clients will not be able to access the file system. This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-acl {enabled|disabled}] - NFSv4.1 ACL Support

This optional parameter specifies whether Data ONTAP supports NFSv4.1 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` . This parameter is not supported for Vservers with Infinite Volume.

[-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 Data ONTAP supports NFSv4.1 read delegations. The default setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-write-delegation {enabled|disabled}] - NFSv4.1 Write Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4.1 write delegations. The default setting is `disabled` . This parameter is not supported for Vservers with Infinite Volume.

[-v4.x-session-num-slots <integer>] - Number of Slots in the NFSv4.x Session slot tables (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) (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 (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 (privilege: advanced)

This optional parameter specifies whether Data ONTAP supports fetching auxillary 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 (privilege: advanced)

This optional parameter specifies the maximum number of auxillary groups supported over RPC security flavors AUTH_SYS and RPCSEC_GSS in Data 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 (privilege: advanced)

This optional parameter specifies whether clustered Data 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 (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 (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 (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 (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 `des` ,`des3` ,`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`* .

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 (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 (privilege: advanced)

If you enable this optional parameter, during client access check evaluation in a netgroup, Data 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 (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 (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*. This parameter is not supported for Vservers with Infinite Volume.

[-ignore-nt-acl-for-root {enabled|disabled}] - Ignore the NT ACL Check for NFS User 'root' (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 (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 from 60000 through 604800000. The default setting is `86400000`.

[-cached-cred-negative-ttl <integer>] - Time To Live Value (in msecs) of a Negative Cached Credential (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 from 60000 through 604800000. The default setting is `7200000`.

[-skip-root-owner-write-perm-check {enabled|disabled}] - Skip Permission Check for NFS Write Calls from Root/Owner (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 (privilege: advanced)

This optional parameter specifies whether Data 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 (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 (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 (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) (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 (privilege: advanced)

This optional parameter specifies whether to enable replay cache checksum for NFS requests . The default value is *enabled* .

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

- vserver export-policy rule create

vserver nfs delete

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.

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
```

vserver nfs modify

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 (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 (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. This parameter is not supported for Vservers with Infinite Volume.

`[-udp {enabled|disabled}]` - UDP Protocol

This optional parameter specifies whether to enable NFS access over UDP. This value is not modifiable on a Vserver with Infinite Volume.

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 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 (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 (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 (privilege: advanced)

This optional parameter specifies whether Data 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 (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 (privilege: advanced)

This optional parameter specifies how NFSv3 security changes affect NTFS volumes. If you set this parameter to `ignore` , Data 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` , Data 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 (privilege: advanced)

This optional parameter specifies whether ownership of a file can be changed by superusers or by non-root users who currently own the file. If you set this parameter to `restricted` , the ownership of a file can be changed by superusers only. If you set this parameter to `unrestricted` , the ownership of a file can be changed by superusers and the current owner of the file. If you set this parameter to `use-export-policy` , the ownership of a file can be changed in accordance with the relevant export rules.

[-trace-enabled {true|false}] - NFS Response Trace Enabled (privilege: advanced)

This optional parameter specifies whether Data 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) (privilege: advanced)

This optional parameter specifies the amount of time, in seconds, after which Data 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) (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) (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.

[-v3-tcp-max-read-size <integer>] - NFSv3 TCP Maximum Read Size (bytes) (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 read requests. The range is 8192 to 1048576. The default setting is 65536 when created.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-v3-tcp-max-write-size <integer>] - NFSv3 TCP Maximum Write Size (bytes) (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 write requests. The range is 8192 to 65536. The default setting is 65536 when created.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-v4.0-acl {enabled|disabled}] - NFSv4.0 ACL Support

This optional parameter specifies whether Data ONTAP supports NFSv4.0 access control lists (ACLs). The default setting is `disabled` when created. This parameter is not supported for Vservers with Infinite Volume.

[-v4.0-read-delegation {enabled|disabled}] - NFSv4.0 Read Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4 read delegations. The default setting is `disabled` when created. This parameter is not supported for Vservers with Infinite Volume.

[-v4.0-write-delegation {enabled|disabled}] - NFSv4.0 Write Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4 write delegations. The default setting is `disabled` when created. This parameter is not supported for Vservers with Infinite Volume.

[-v4-fsid-change {enabled|disabled}] - Show Change in FSID as NFSv4 Clients Traverse Filesystems (privilege: advanced)

This optional parameter specifies whether Data 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.

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 (privilege: advanced)

This optional parameter specifies whether Data 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. This parameter is not supported for Vservers with Infinite Volume.

[-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 normally taken from the NIS domain or the DNS domain in use. 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 (privilege: advanced)

This optional parameter specifies whether Data 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) (privilege: advanced)

This optional parameters specifies the time period in which Data 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 Data 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 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. This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-implementation-domain <nfs domain>] - NFSv4.1 Implementation ID Domain (privilege: advanced)

This optional parameter specifies the NFSv4.1 implementation domain.

[-v4.1-implementation-name <text>] - NFSv4.1 Implementation ID Name (privilege: advanced)

This optional parameter specifies the NFSv4.1 implementation name.

[-v4.1-implementation-date <Date>] - NFSv4.1 Implementation ID Date (privilege: advanced)

This optional parameter specifies the NFSv4.1 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. The default setting is `enabled` at the time of creation.

[-v4.1-referrals {enabled|disabled}] - NFSv4.1 Referral Support (privilege: advanced)

This optional parameter specifies whether Data ONTAP supports NFSv4.1 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 referrals, set this option to `disabled` ; otherwise, those clients will not be able to access the file system. This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-acl {enabled|disabled}] - NFSv4.1 ACL Support

This optional parameter specifies whether Data ONTAP supports NFSv4.1 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. This parameter is not supported for Vservers with Infinite Volume.

[-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 Data ONTAP supports NFSv4.1 read delegations. The default setting is `disabled` when created. This parameter is not supported for Vservers with Infinite Volume.

[-v4.1-write-delegation {enabled|disabled}] - NFSv4.1 Write Delegation Support

This optional parameter specifies whether Data ONTAP supports NFSv4.1 write delegations. The default setting is `disabled` when created. This parameter is not supported for Vservers with Infinite Volume.

[-v4.x-session-num-slots <integer>] - Number of Slots in the NFSv4.x Session slot tables (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) (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 (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 (privilege: advanced)

This optional parameter specifies whether Data ONTAP supports fetching auxillary 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 (privilege: advanced)

This optional parameter specifies the maximum number of auxillary groups supported over RPC security flavors AUTH_SYS and RPCSEC_GSS in Data 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 (privilege: advanced)

This optional parameter specifies whether clustered Data 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 (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 (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 (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 (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 `des` ,`des3` ,`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`* .

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 (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 (privilege: advanced)

If you enable this optional parameter, during client access check evaluation in a netgroup, Data 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 (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 (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. This parameter is not supported for Vservers with Infinite Volume

[-ignore-nt-acl-for-root {enabled|disabled}] - Ignore the NT ACL Check for NFS User 'root' (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 (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 from 60000 through 604800000. The default setting is `86400000`.

[-cached-cred-negative-ttl <integer>] - Time To Live Value (in msecs) of a Negative Cached Credential (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 from 60000 through 604800000. The default setting is `7200000`.

[-skip-root-owner-write-perm-check {enabled|disabled}] - Skip Permission Check for NFS Write Calls from Root/Owner (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 (privilege: advanced)

This optional parameter specifies whether Data 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 (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 (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 (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) (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 (privilege: advanced)

This optional parameter specifies whether to enable replay cache checksum for NFS requests. The default value is *enabled* .

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

- vserver export-policy rule create

vserver nfs off

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
```

vserver nfs on

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
```

vserver nfs prepare-for-v3-ms-dos-client-downgrade

Disable NFSv3 MS-DOS Client Support

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

Description

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 Data 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
```

vserver nfs prepare-to-downgrade

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 Data ONTAP.

Parameters

-disable-feature-set <downgrade version> - Data ONTAP Version (privilege: advanced)

This parameter specifies the Data 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 Data 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
```

vserver nfs show

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.

Parameters

{ [-fields <fieldname>,…]

If you specify the `-fields` parameter, the command only displays the fields that you specify.

| [-krb-opts] (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 (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 (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 (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 (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 (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 (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 (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 (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 (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) (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) (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) (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.

[-v3-tcp-max-read-size <integer>] - NFSv3 TCP Maximum Read Size (bytes) (privilege: advanced)

If you specify this parameter, the command displays information only about the NFS-enabled Vservers with the specified TCP maximum transfer size for NFSv3 read requests. The range is 8192 to 1048576.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-v3-tcp-max-write-size <integer>] - NFSv3 TCP Maximum Write Size (bytes) (privilege: advanced)

If you specify this parameter, the command displays information only about the NFS-enabled Vservers with the specified TCP maximum transfer size for NFSv3 write requests. The range is 8192 to 65536.

This parameter is deprecated and may be removed in a future release of Data ONTAP. Use the `-tcp-max-xfer-size` parameter instead.

[-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 (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 (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 (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) (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 (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 (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 (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 (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 (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) (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 (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 (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 (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 (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 (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 (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 (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 (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 : des, des3, 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 (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 (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 (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 (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' (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 (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 (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 (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 (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 (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 (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 (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) (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 (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.

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.
```

vserver nfs start

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
```

vserver nfs status

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.
```

vserver nfs stop

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
```

vserver nfs kerberos interface disable

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
```

vserver nfs kerberos interface enable

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 {(ftp|http)://(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".

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
```

vserver nfs kerberos interface modify

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 <lif-name> - 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 {(ftp|http)://(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".

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
```

vserver nfs kerberos interface show

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 <lif-name>] - 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.

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
```

vserver nfs kerberos realm create

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 optional 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 seconds of clock skew between the clients and the server are permitted. The default setting is 300 seconds.

`[-adserver-name <text>]` - Active Directory Server Name

This optional 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 optional 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
```

vserver nfs kerberos realm delete

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
```

vserver nfs kerberos realm modify

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 optional 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 seconds of clock-skew between server and the clients are permitted. The default setting at the time of creation is 300 seconds.

[-adserver-name <text>] - Active Directory Server Name

This optional 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 optional 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
```

vserver nfs kerberos realm show

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
```

vserver nfs pnfs devices create

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 Data 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 (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 (privilege: advanced)

This parameter specifies the Vserver to which the volumes belong.

`-msid <integer>` - Volume MSID (privilege: advanced)

This parameter uniquely identifies the volume for which you are creating a pNFS device.

`-striping-epoch <integer>` - Striping Epoch (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 (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 (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 (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 (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 (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

vserver nfs pnfs devices delete

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 vserver nfs pnfs devices show .

Parameters

`-global-device-table-id <integer>` - Global Device Mapping Table ID (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

- vserver nfs pnfs devices show

vserver nfs pnfs devices show

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] }

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

[-global-device-table-id <integer>] - Global Device Mapping Table ID (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 (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 (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 (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 (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 (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 (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 (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 (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
```

vserver nfs pnfs devices cache show

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] }

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

[-node {<nodename>|local}] - Node (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 (privilege: advanced)

If you specify this parameter, the command displays information only about the Vserver that has the pNFS devices cache.

Examples

vserver nfs pnfs devices mappings show

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] }

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

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

If your specify this parameter, the command displays information only about the Vserver that the mapping identifier and DSID belong to.

[-global-device-table-id <integer>] - Global Device Mapping Table ID (privilege: advanced)

This specifies the unique identifier that the pNFS devices subsystem assigns to the device whose mappings are being output.

[-dsid <integer>] - Constituent Volume DSID (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 (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

- vserver nfs pnfs devices show

vserver peer commands

vserver peer accept

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 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

- vserver peer create

vserver peer create

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}` - 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
```

vserver peer delete

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.

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
```

vserver peer modify-local-name

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
```

vserver peer modify

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}` - 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
```

vserver peer reject

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
```

vserver peer repair-peer-name

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 (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
```

vserver peer resume

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
```

vserver peer show-all

(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 (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 (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}] - 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.
```

vserver peer show

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}] - 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 (privilege: advanced)

If this parameter is specified, the command displays relationships that match the specified peer Vserver UUID.

[-vserver-uuid <UUID>] - Local Vserver UUID (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.
```

vserver peer suspend

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
```

vserver peer permission create

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 explict 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>,... - 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

- vserver peer accept

vserver peer permission delete

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
```

vserver peer permission show

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>,…] - 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.
```

vserver peer transition create

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.

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 -vserver vs1.example.com -src
-filer-name src1.example.com -multi-path-address src1-e0d.example.com
-local-lifs lif1,lif2
```

vserver peer transition delete

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 -vserver lvs1.example.com -src
-filer-name src1.example.com
```

vserver peer transition modify

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 -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 -vserver vs1.example.com -src
-filer-name src1.example.com
-local-lifs lif1,lif2
```

vserver peer transition show

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
```

vserver san commands

vserver san prepare-to-downgrade

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 Data 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 DATA 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 DATA ONTAP release.

vserver security commands

vserver security file-directory apply

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.

> 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). Amongst other things, the task defines which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.

The `vserver security file-directory apply` command is not supported for Vservers with Infinite Volume.

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 (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
```

vserver security file-directory remove-slag

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.

The `vserver security file-directory remove-slag` command is not supported for Vservers with Infinite Volume.

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: -
```

vserver security file-directory show-effective-permissions

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 con-tained 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.

vserver security file-directory show

Display file/folder security information

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

Description

The `vserver 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.

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.

[-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)@Devic
e.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
```

vserver security file-directory job show

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.

| [-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.

| [-jobuuid] (privilege: advanced)

Displays the job ID, the job name, the owning Vserver, and the job UUID.

| [-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

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 abthe 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
```

vserver security file-directory ntfs create

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.

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). Amongst other things, the task defines which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.

The `vserver security file-directory ntfs create` command is not supported for Vservers with Infinite Volume.

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.

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

> 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 (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

> 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 (privilege: advanced)

Specifies the control flags in the security descriptor.

> 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
```

vserver security file-directory ntfs delete

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).

The `vserver security file-directory ntfs delete` command is not supported for Vservers with Infinite Volume.

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
```

vserver security file-directory ntfs modify

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.

The `vserver security file-directory ntfs modify` command is not supported for Vservers with Infinite Volume.

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

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 (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

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 (privilege: advanced)

Specifies the control flags in the security descriptor to be modified.

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
```

vserver security file-directory ntfs show

Display an NTFS security descriptors

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

Description

The `vserver 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.

The `vserver security file-directory ntfs show` command is not supported for Vservers with Infinite Volume.

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 (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 (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
```

vserver security file-directory ntfs dacl add

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.

> 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). Amongst other things, the task defines which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.

The `vserver security file-directory ntfs dacl add` command is not supported for Vservers with Infinite Volume.

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

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 (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

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
```

vserver security file-directory ntfs dacl modify

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

The `vserver security file-directory ntfs dacl modify` command is not supported for Vservers with Infinite Volume.

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

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 (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

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
```

vserver security file-directory ntfs dacl remove

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

The `vserver security file-directory ntfs dacl remove` command is not supported for Vservers with Infinite Volume.

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
```

vserver security file-directory ntfs dacl show

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.

The `vserver security file-directory ntfs dacl show` command is not supported for Vservers with Infinite Volume.

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

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 (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.
```

vserver security file-directory ntfs sacl add

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.

> 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). Amongst other things, the task defines which security descriptor to apply to a path.

* Apply a policy to the associated Vserver.

The `vserver security file-directory ntfs sacl add` command is not supported for Vservers with Infinite Volume.

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

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

| `[-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

| [-rights-raw <Hex Integer>] - Raw Access Rights (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 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

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
```

vserver security file-directory ntfs sacl modify

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

The `vserver security file-directory ntfs sacl modify` command is not supported for Vservers with Infinite Volume.

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

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 (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 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

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
```

vserver security file-directory ntfs sacl remove

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

The `vserver security file-directory ntfs sacl remove` command is not supported for Vservers with Infinite Volume.

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
```

vserver security file-directory ntfs sacl show

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.

The `vserver security file-directory ntfs sacl show` command is not supported for Vservers with Infinite Volume.

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

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 (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
```

vserver security file-directory policy create

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.

> 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.

> 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). Amongst other things, the task defines which security descriptor to apply to a path.
* Apply a policy to the associated Vserver.

The `vserver security file-directory policy create` command is not supported for Vservers with Infinite Volume.

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
```

vserver security file-directory policy delete

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.
```

The `vserver security file-directory policy delete` command is not supported for Vservers with Infinite Volume.

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
```

vserver security file-directory policy show

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.

The `vserver security file-directory policy show` command is not supported for Vservers with Infinite Volume.

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.
```

vserver security file-directory policy task add

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.

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.

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). Amongst other things, the task defines which security descriptor to apply to a path.

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.

The `vserver security file-directory policy task add` command is not supported for Vservers with Infinite Volume.

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.

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”.

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

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 FlexGroups. 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
```

vserver security file-directory policy task modify

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.

> 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

> The only security type supported in this Data ONTAP release is *"ntfs"*; therefore, you cannot modify the `-security-type` parameter.

The `vserver security file-directory policy task modify` command is not supported for Vservers with Infinite Volume.

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.

> 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*.

> 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
```

vserver security file-directory policy task remove

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.

The `vserver security file-directory policy task remove` command is not supported for Vservers with Infinite Volume.

> 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 /
```

vserver security file-directory policy task show

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

The `vserver security file-directory policy task show` command is not supported for Vservers with Infinite Volume.

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.

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.
```

vserver security trace filter create

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 Data ONTAP 9.3, this feature was only supported for CIFS. In Data ONTAP 9.3 and later, this feature is supported for both NFS and CIFS.

The `vserver security trace filter create` command is not supported for Vservers with Infinite Volume.

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 disabled.

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. 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
```

vserver security trace filter delete

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 Data ONTAP 9.3, this feature was only supported for CIFS. In Data ONTAP 9.3 and later, this feature is supported for both NFS and CIFS.

The `vserver security trace filter delete` command is not supported for Vservers with Infinite Volume.

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 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

- vserver security trace filter show

vserver security trace filter modify

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 Data ONTAP 9.3, this feature was only supported for CIFS. In Data ONTAP 9.3 and later, this feature is supported for both NFS and CIFS.

The `vserver security trace filter modify` command is not supported for Vservers with Infinite Volume.

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 disabled.

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
```

vserver security trace filter show

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 Data ONTAP 9.3, this feature was only supported for CIFS. In Data ONTAP 9.3 and later, this feature is supported for both NFS and CIFS.

The `vserver security trace filter show` command is not supported for Vservers with Infinite Volume.

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 specififc 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 only about filters that are disabled after the specified minutes.

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.
```

vserver security trace trace-result delete

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.
```

The `vserver security trace trace-result delete` command is not supported for Vservers with Infinite Volume.

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
```

vserver security trace trace-result show

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 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_`` .
```

The `vserver security trace trace-result show` command is not supported for Vservers with Infinite Volume.

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

- vserver security trace filter create

vserver services commands

vserver services access-check authentication get-claim-name

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 (privilege: advanced)

The name of the node on which the command is executed.

-vserver <vserver> - Vserver Name (privilege: advanced)

The name of the Vserver.

-claim-cn <text> - Claim CN (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
```

vserver services access-check authentication get-dc-info

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 (privilege: advanced)

The name of the node on which the command is executed.

-vserver <vserver> - Vserver Name (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
```

vserver services access-check authentication sid-to-uid

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 (privilege: advanced)

The name of the node on which the command is executed.

`-vserver <vserver>` - Vserver Name (privilege: advanced)

The name of the Vserver.

`-sid <text>` - Windows SID (privilege: advanced)

The SID of a Windows user.

`[-clientIp <IP Address>]` - Client IP Address (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
```

vserver services access-check authentication sid-to-unix-name

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 (privilege: advanced)

The name of the node on which the command is executed.

`-vserver <vserver>` - Vserver Name (privilege: advanced)

The name of the Vserver.

`-sid <text>` - Windows SID (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
```

vserver services access-check authentication translate

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 (privilege: advanced)

The name of the node on which the command is executed.

-vserver <vserver> - Vserver Name (privilege: advanced)

The name of the Vserver.

{ -uid <integer> - UNIX User ID (privilege: advanced)

The UNIX user's UID.

| -gid <integer> - UNIX Group ID (privilege: advanced)

The UNIX user's GID.

| -sid <text> - Windows SID (privilege: advanced)

The Windows user's SID.

| -unix-user-name <text> - UNIX User Name (privilege: advanced)

The UNIX username.

| -unix-group-name <text> - UNIX Group Name (privilege: advanced)

The UNIX group name.

| -win-name <text> - Windows Name (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
```

vserver services access-check authentication uid-to-sid

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 (privilege: advanced)

The name of the node on which the command is executed.

`-vserver <vserver>` - Vserver Name (privilege: advanced)

The name of the Vserver.

`-uid <integer>` - UNIX User ID (privilege: advanced)

The User ID of a UNIX user.

`[-clientIp <IP Address>]` - Client IP Address (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
```

vserver services name-service cache group-membership delete-all

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 group-membership delete command removes the cached group membership entries of the users for the specified Vserver.

Parameters

`-vserver <vserver name>` - Vserver (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

- vserver services name-service cache group-membership delete

vserver services name-service cache group-membership delete

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 (privilege: advanced)

Use this parameter to specify the Vserver for which the group membership entries need to be deleted.

`-user <text>` - User Name (privilege: advanced)

Use this parameter to specify the user name for which the cached group membership entries need to be deleted.

`-group <integer>` - Gid (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
```

vserver services name-service cache group-membership show

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] }

Use this parameter to display detailed information about the cached group membership details of the user.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the group membership entries need to be displayed.

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

Use this parameter to display information only about the cached group membership entries that have the specified user name.

[-group <integer>] - Gid (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 (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 (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 (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 (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
```

vserver services name-service cache group-membership settings modify

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 (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? (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 diabled. The default value for this parameter is *`true`* .

`[-grplist-ttl <[<integer>h][<integer>m][<integer>s]>]` - Time to Live for Grplist (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 24 hours.

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 -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
```

vserver services name-service cache group-membership settings show

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] }

Use this parameter to display detailed information about the group membership cache configuration settings.

[-vserver <vserver name>] - Vserver (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? (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>h][<integer>m][<integer>s]>] - Time to Live for Grplist (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
```

vserver services name-service cache hosts forward-lookup delete-all

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 (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
```

vserver services name-service cache hosts forward-lookup delete

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 (privilege: advanced)

Use this parameter to specify the Vserver for which the cached forward lookup table entries need to be deleted.

`-host <text>` - Hostname (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 (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 (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 (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 (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
```

vserver services name-service cache hosts forward-lookup show

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] }

Use this parameter to display detailed information about the cached forward lookup table entries.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the cached forward lookup table entries need to be displayed.

[-host <text>] - Hostname (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 (privilege: advanced)

Use this parameter to display information only about the cached forward lookup table entries that have the specifed protocol.

[-sock-type {SOCK_ANY|SOCK_STREAM|SOCK_DGRAM|SOCK_RAW}] - Sock Type (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 (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 (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 (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 (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 (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 (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 (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 (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}] - Source of the Entry (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 (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 (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
```

vserver services name-service cache hosts reverse-lookup delete-all

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 (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
```

vserver services name-service cache hosts reverse-lookup delete

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 (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 (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 (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
```

vserver services name-service cache hosts reverse-lookup show

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] }

Use this parameter to display detailed information about the cached reverse lookup table entries.

[-vserver <vserver name>] - Vserver (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 (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 (privilege: advanced)

Use this parameter to display information only about the cached reverse lookup table entries that have the specifed service flag.

[-host <text>] - Hostname (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 (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 (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 (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 (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 (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}] - Source of the Entry (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 (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
```

vserver services name-service cache hosts settings modify

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 (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? (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? (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>h][<integer>m][<integer>s]>]` - Time to Live (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>h][<integer>m][<integer>s]>]` - Negative Time to Live (privilege: advanced)

Use this parameter to specify the time for which the negative entries need to be cached. The default value is 1 minute.

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
```

vserver services name-service cache hosts settings show

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] }

Use this parameter to display detailed information about the hosts cache configuration settings.

[-vserver <vserver name>] - Vserver (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? (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? (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>h][<integer>m][<integer>s]>] - Time to Live (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>h][<integer>m][<integer>s]>] - Negative Time to Live (privilege: advanced)

Use this parameter to display information only about the hosts cache configuration setings that have the specified negative Time to Live.

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
```

vserver services name-service cache netgroups ip-to-netgroup delete-all

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 (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
```

vserver services name-service cache netgroups ip-to-netgroup delete

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 (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 (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 (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
```

vserver services name-service cache netgroups ip-to-netgroup show

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] }

Use this parameter to display detailed information about the cached client IP to netgroup entries.

[-vserver <vserver name>] - Vserver (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 (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 (privilege: advanced)

Use this parameter to display information only about the cached IP to netgroup entries that have the specifed netgroup.

[-create-time <MM/DD/YYYY HH:MM:SS>] - Create Time (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}] - Source of the Entry (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 deletes all the cached IP to netgroup entries for Vserver vs0:

```
cluster1::> vserver services name-service cache netgroups ip-to-netgroup
show -vserver vs0
```

vserver services name-service cache netgroups members delete-all

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 (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
```

vserver services name-service cache netgroups members delete

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 (privilege: advanced)

Use this parameter to specify the Vserver for which the cached netgroup members entries need to be deleted.

`-netgroup <text>` - Netgroup (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
```

vserver services name-service cache netgroups members show

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 ]` }

Use this parameter to display detailed information about the cached members of a netgroup.

`[-vserver <vserver name>]` - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the cached netgroup members entries need to be displayed.

`[-netgroup <text>]` - Netgroup (privilege: advanced)

Use this parameter to display information only about the cached members that belong to the specified netgroup.

`[-hosts <text>]` - Hosts (privilege: advanced)

Use this parameter to display information only about the cached netgroups that have the specifed host as a member.

[-create-time <MM/DD/YYYY HH:MM:SS>] - Create Time (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}] - Source of the Entry (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
```

vserver services name-service cache netgroups settings modify

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 (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? (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? (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>h][<integer>m][<integer>s]>] - Time to Live (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>h][<integer>m][<integer>s]>] - Negative Time to Live (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>h][<integer>m][<integer>s]>] - TTL for netgroup members (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
```

vserver services name-service cache netgroups settings show

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] }

Use this parameter to display detailed information about the netgroups cache configuration settings.

[-vserver <vserver name>] - Vserver (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? (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? (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>h][<integer>m][<integer>s]>] - Time to Live (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>h][<integer>m][<integer>s]>] - Negative Time to Live (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>h][<integer>m][<integer>s]>] - TTL for netgroup members (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
```

vserver services name-service cache unix-group group-by-gid delete-all

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 (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
```

vserver services name-service cache unix-group group-by-gid delete

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 (privilege: advanced)

Use this parameter to specify the Vserver for which the group entries that are cached by the group identifer or GID need to be deleted.

`-gr-gid <integer>` - gr_gid field (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
```

vserver services name-service cache unix-group group-by-gid show

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] }

Use this parameter to display detailed information about the group entries cached by the group identifier or GID.

[-vserver <vserver name>] - Vserver (privilege: advanced)

Use this parameter to specify the Vserver for which the group entries that are cached by the group identifer or GID need to be displayed.

[-gr-gid <integer>] - gr_gid field (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 (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 (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}] - Source of the Entry (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
```

vserver services name-service cache unix-group group-by-name delete-all

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 (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
```

vserver services name-service cache unix-group group-by-name delete

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 (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 (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
```

vserver services name-service cache unix-group group-by-name show

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] }

Use this parameter to display detailed information about the group entries cached by group name.

[-vserver <vserver name>] - Vserver (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 (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 (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 (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}] - Source of the Entry (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
```

vserver services name-service cache unix-group settings modify

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 (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? (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? (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>h][<integer>m][<integer>s]>]` - Time to Live (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>h][<integer>m][<integer>s]>]` - Negative Time to Live (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? (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
```

vserver services name-service cache unix-group settings show

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] }

Use this parameter to display detailed information about the groups cache configuration settings.

[-vserver <vserver name>] - Vserver (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? (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? (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>h][<integer>m][<integer>s]>] - Time to Live (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>h][<integer>m][<integer>s]>] - Negative Time to Live (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? (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
```

vserver services name-service cache unix-user settings modify

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 (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? (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? (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>h][<integer>m][<integer>s]>]` - Time to Live (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>h][<integer>m][<integer>s]>]` - Negative Time to Live (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? (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
```

vserver services name-service cache unix-user settings show

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] }

Use this parameter to display detailed information about the users cache configuration settings.

[-vserver <vserver name>] - Vserver (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? (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? (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>h][<integer>m][<integer>s]>] - Time to Live (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>h][<integer>m][<integer>s]>] - Negative Time to Live (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? (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
```

vserver services name-service cache unix-user user-by-id delete-all

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 (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
```

vserver services name-service cache unix-user user-by-id delete

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 (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 (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 identifer 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
```

vserver services name-service cache unix-user user-by-id show

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] }

Use this parameter to display detailed information about the user entries cached by the user identifier or UID.

[-vserver <vserver name>] - Vserver (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 (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 (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 (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 (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}] - Source of the Entry (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
```

vserver services name-service cache unix-user user-by-name delete-all

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 (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
```

vserver services name-service cache unix-user user-by-name delete

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 (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 (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
```

vserver services name-service cache unix-user user-by-name show

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 ]` }

Use this parameter to display detailed information about the user entries cached by the user name.

[-vserver <vserver name>] - Vserver (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 (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 (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 (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 (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}] - Source of the Entry (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
```

vserver services name-service dns check

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.
```

vserver services name-service dns create

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? (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 (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 (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 254 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
```

vserver services name-service dns delete

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
```

vserver services name-service dns modify

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 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? (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 (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 (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 254 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

- vserver services name-service dns delete

vserver services name-service dns show

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? (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 (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 (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
```

vserver services name-service dns dynamic-update modify

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>h][<integer>m][<integer>s]>]` - Time to Live for DNS Updates (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
```

vserver services name-service dns dynamic-update prepare-to-downgrade

Disable the Dynamic DNS update feature to be compatible with releases earlier than Data ONTAP 8.3.1

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

Description

The `vserver services name-service dns dynamic-update prepare-to-downgrade` command disables the Dynamic DNS updates on all Vservers and removes all related configurations. This command is used to prepare for downgrading the system to a release earlier than Data ONTAP 8.3.1 only.

Examples

The following example disables the dynamic DNS updates feature.

```
cluster1::*> vserver services name-service dns dynamic-update prepare-to-
downgrade
  Warning: This command will disable dynamic DNS updates on all Vservers,
remove
    all related configurations, and disable the dynamic DNS update
    feature. Use this command to prepare for downgrading the system to a
    release earlier than Data ONTAP 8.3.1 only.
  Do you want to continue? {y|n}:
```

vserver services name-service dns dynamic-update show

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>h][<integer>m][<integer>s]>] - Time to Live for DNS Updates (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.
```

vserver services name-service dns dynamic-update record add

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 (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 (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
```

vserver services name-service dns dynamic-update record delete

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 (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 (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 (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
```

vserver services name-service dns hosts create

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
```

vserver services name-service dns hosts delete

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.
```

vserver services name-service dns hosts modify

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.
```

vserver services name-service dns hosts show

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.
```

vserver services name-service getxxbyyy getaddrinfo

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-hostname <text>` - Host Name (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 (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 (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Enable/Disable cache (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
```

vserver services name-service getxxbyyy getgrbygid

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-groupID <integer>` - Group ID (privilege: advanced)

Use this parameter to specify the GroupID for which the members are requested

`[-show-source {true|false}]` - Source used for Lookup (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Use Locally-Cached Values (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
```

vserver services name-service getxxbyyy getgrbyname

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 (privilege: advanced)
Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)
Use this parameter to specify the Vserver where the lookup will be performed

`-groupname <text>` - Group Name (privilege: advanced)
Use this parameter to specify the Group Name for which the members are requested

`[-show-source {true|false}]` - Source used for Lookup (privilege: advanced)
Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Use Locally-Cached Values (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
```

vserver services name-service getxxbyyy getgrlist

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

-vserver <vserver name> - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

-username <text> - User Name (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 (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
```

vserver services name-service getxxbyyy gethostbyaddr

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

-vserver <vserver name> - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

-ipaddress <IP Address> - IP Address (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 (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

[-use-cache <true>] - Enable/Disable cache (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
```

vserver services name-service getxxbyyy gethostbyname

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 (privilege: advanced)

Node Use this parameter to specify the node where the lookup will be performed

-vserver <vserver name> - Vserver Name (privilege: advanced)

Vserver Name Use this parameter to specify the Vserver where the lookup will be performed

-hostname <text> - Host Name (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 (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
```

vserver services name-service getxxbyyy getnameinfo

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-ipaddress <IP Address>` - IP Address (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 (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Enable/Disable cache (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
```

vserver services name-service getxxbyyy getpwbyname

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-username <text>` - User Name (privilege: advanced)

Use this parameter to specify the Username for which the password entry is requested

`[-show-source {true|false}]` - Source used for Lookup (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Enable/Disable cache (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
```

vserver services name-service getxxbyyy getpwbyuid

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-userID <integer>` - User ID (privilege: advanced)

Use this parameter to specify the UserID for whom the password entry is requested

`[-show-source {true|false}]` - Source used for Lookup (privilege: advanced)

Use this parameter to specify if source used for lookup needs to be displayed

`[-use-cache <true>]` - Enable/Disable cache (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: $1$f7b22f68$KihT1ptYqpEjcM4jfE60f0
      pw_uid: 1001
      pw_gid: 65533
      pw_gecos: User
      pw_dir: /var/home/vsadmin
      pw_shell: /sbin/ngsh
```

vserver services name-service getxxbyyy netgrp

Checks if a client is part of a netgroup.

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

Description

The ` vserver services name-service getxxbyyy netgrp ` 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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-netgroup <text>` - Netgroup Name (privilege: advanced)

Use this parameter to specify the Netgroup name

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

Use this parameter to specify the Client name for which the membership in a given netgroup needs to be checked

`[-show-source {true|false}]` - Source used for Lookup (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 netgrp -node
cluster1-01 -vserver vs1 -netgroup net1 -client-name h1 -show-source false
    h1 is a member of net1
```

vserver services name-service getxxbyyy netgrpbyhost

Check if a client is part of a netgroup using netgroup-by-host query

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

Description

The ` vserver services name-service getxxbyyy netgrpbyhost` command checks whether a client is part of a netgroup using the netgroup.byhost map. The underlying service for doing the lookup is selected based on the configured name service switch order.

Parameters

`-node {<nodename>|local}` - Node Name (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed.

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed.

`-netgroup <text>` - Netgroup Name (privilege: advanced)

Use this parameter to specify the netgroup name.

`-clientIP <IP Address>` - Client IP Address (privilege: advanced)

Use this parameter to specify the client IPv4/IPv6 address for which you want to check the membership in a given netgroup.

`[-enable-domain-search-flag {true|false}]` - Use DNS domain (privilege: advanced)

Use this parameter to specify whether you want to perform shortname host lookups in case the configured DNS search domains match the domain returned by the reverse lookup.

`[-show-source {true|false}]` - Source Used for Lookup (privilege: advanced)

Use this parameter to specify whether you want to display the source used for the lookup.

Examples

The following example checks whether the given client IP address is a member of the given netgroup:

```
cluster1::*> vserver services name-service getxxbyyy netgrpbyhost -node
node1 -vserver vs1 -netgroup ng1 -clientIP 10.10.10.90
    Hostname resolved to: dnshost.example.com
    Success
```

vserver services name-service getxxbyyy netgrpcheck

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 (privilege: advanced)

Use this parameter to specify the node where the lookup will be performed

`-vserver <vserver name>` - Vserver Name (privilege: advanced)

Use this parameter to specify the Vserver where the lookup will be performed

`-netgroup <text>` - Netgroup Name (privilege: advanced)

Use this parameter to specify the Netgroup name

`-clientIP <IP Address>` - Client IP Address (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 (privilege: advanced)

Use this parameter to use DNS domain. Default value for this field is true

`[-trust-any-source {true|false}]` - Trust any source (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 (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
```

vserver services name-service ldap check

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.

Examples

The following example checks the LDAP configuration on the Vserver 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".
```

vserver services name-service ldap create

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
```

vserver services name-service ldap delete

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.

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
```

vserver services name-service ldap modify

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.

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
```

vserver services name-service ldap show

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
-------------- -------------
vs1            corp
```

vserver services name-service ldap client create

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` , `-preferred-ad-servers` or `-bind-as-cifs-server` 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 onwards. 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.

`[-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 the -ad-domain is configured, the default is `true` , otherwise the default is `false` .

`-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 that the LDAP client uses to connect to LDAP servers. If you do not specify this parameter, the default is port `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 an `anonymous` bind.

[-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 (CN=user,DC=domain,DC=com) form. 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 (privilege: advanced)

This parameter specifies the user DN, which overrides the base DN for user lookups.

> 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 (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 (privilege: advanced)

This parameter specifies the group DN, which overrides the base DN for group lookups.

> 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 (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 (privilege: advanced)

This parameter specifies the netgroup DN, which overrides the base DN netgroup lookups.

> 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 (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 (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 (privilege: advanced)

This parameter specifies the netgroup-by-host DN, which overrides the base DN for netgroup-by-host lookups.

> 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 (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.

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.

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
```

vserver services name-service ldap client delete

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
```

vserver services name-service ldap client modify-bind-password

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:
```

vserver services name-service ldap client modify

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.

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` , `-preferred-ad-servers` or `-bind-as-cifs-server` 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 onwards. 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.

[-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. If you do not specify this parameter, the default is `false` .

[-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 that the LDAP client uses to connect to LDAP servers. If you do not specify this parameter, the default is port `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 an `anonymous` bind.

[-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 (CN=user,DC=domain,DC=com) form. 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 (privilege: advanced)

This parameter specifies the user DN, which overrides the base DN for user lookups.

> 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 (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 (privilege: advanced)

This parameter specifies the group DN, which overrides the base DN for group lookups.

> 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 (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 (privilege: advanced)

This parameter specifies the netgroup DN, which overrides the base DN netgroup lookups.

> 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 (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 (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 (privilege: advanced)

This parameter specifies the netgroup-by-host DN, which overrides the base DN for netgroup-by-host lookups.

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 (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` .

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
```

vserver services name-service ldap client show

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.

[-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 (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 (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 (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 (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 (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 (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 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 (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 (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 (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.

Examples

The following example shows a summary of all of the LDAP client configurations available for Vserver `vs1` :

```
cluster1::> vserver services name-service ldap 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
```

vserver services name-service ldap client schema copy

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 (privilege: advanced)

This parameter specifies the Vserver for which you want to copy an existing LDAP schema template.

-schema <text> - Schema Template (privilege: advanced)

This parameter specifies the name of the existing schema template you want to copy.

-new-schema-name <text> - New Schema Template Name (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
```

vserver services name-service ldap client schema delete

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.

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
```

vserver services name-service ldap client schema modify

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 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

- vserver services name-service ldap client schema copy

vserver services name-service ldap client schema show

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.
```

vserver services name-service netgroup load

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 Data 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, Data 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 {(ftp|http)://(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 (privilege: advanced)

If this parameter is specified, the hostname validation is skipped.

`[-skip-file-size-check <true>]` - Skip File Size Check Before Download (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.

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 (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
```

vserver services name-service netgroup status

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you specify this parameter, the command displays the netgroup status only for the specified Vserver.

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

vserver services name-service netgroup file delete

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
```

vserver services name-service netgroup file show

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          '-'
```

vserver services name-service nis-domain create

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 }

ⓘ 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
```

vserver services name-service nis-domain delete

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
```

vserver services name-service nis-domain modify

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 vserver services name-service nis-domain delete command and then create the NIS configuration with new NIS domain using the vserver services name-service nis-domain create command. To permanently remove a configuration, use the 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 }

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

- vserver services name-service nis-domain delete
- vserver services name-service nis-domain create

vserver services name-service nis-domain show-bound

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.
```

vserver services name-service nis-domain show

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

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.
```

vserver services name-service nis-domain group-database build

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 (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
```

vserver services name-service nis-domain group-database status

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 ]` }

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

[-vserver <vserver name>] - Vserver (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 (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 (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 (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 (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
```

vserver services name-service ns-switch create

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
```

vserver services name-service ns-switch delete

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.
```

vserver services name-service ns-switch modify

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
```

vserver services name-service ns-switch show

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.
```

vserver services name-service unix-group adduser

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, Data 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
```

vserver services name-service unix-group create

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, Data 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
```

vserver services name-service unix-group delete

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
```

vserver services name-service unix-group deluser

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
```

vserver services name-service unix-group load-from-uri

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 {(ftp|http)://(hostname|IPv4 Address|'['IPv6 Address']')…} - URI to Load From

Use this parameter to specify the URI from which the command loads group information.

[-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, Data 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
```

vserver services name-service unix-group modify

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
```

vserver services name-service unix-group show

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
```

vserver services name-service unix-group file show

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 (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 (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:
```

vserver services name-service unix-group file status

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you specify this parameter, the command displays the UNIX group status only for the specified Vserver.

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

vserver services name-service unix-group max-limit modify

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 (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
```

vserver services name-service unix-group max-limit show

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 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

- vserver services name-service unix-group max-limit modify

vserver services name-service unix-user create

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, Data 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"
```

vserver services name-service unix-user delete

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
```

vserver services name-service unix-user load-from-uri

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 {(ftp|http)://(hostname|IPv4 Address|'['IPv6 Address']')…} - URI to Load From

This specifies the URI from which user information is to be loaded.

[-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, Data 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
```

vserver services name-service unix-user modify

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"
```

vserver services name-service unix-user show

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
```

vserver services name-service unix-user file show

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 (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 (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::::::
```

vserver services name-service unix-user file status

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you specify this parameter, the command displays the UNIX user status only for the specified Vserver.

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

vserver services name-service unix-user max-limit modify

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 (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
```

vserver services name-service unix-user max-limit show

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 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

- vserver services name-service unix-user max-limit modify

vserver services name-service ypbind start

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
```

vserver services name-service ypbind status

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.

Examples

The following example displays ypbind status:

```
vs1::> vserver services name-service ypbind status
      Status: Running
```

vserver services name-service ypbind stop

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
```

vserver services ndmp generate-password

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
```

vserver services ndmp kill-all

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
```

vserver services ndmp kill

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
```

vserver services ndmp modify

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 262,144(256K), 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 262,144(256K).

[-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 (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 (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 (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 (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 (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 (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 (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 {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 (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::>
```

vserver services ndmp off

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
```

vserver services ndmp on

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
```

vserver services ndmp probe

Display list of NDMP sessions

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

Description

The 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

- system services ndmp probe

vserver services ndmp show

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 (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 (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^moverdata 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 (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 (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 (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 (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 (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 {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 (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::*>
```

vserver services ndmp status

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.
```

vserver services ndmp extensions modify

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 (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 extenion 0x2050 in the Vserver-aware NDMP mode of operation:

```
cluster1::> vserver services ndmp extension modify -is-extension-0x2050
-enabled true
    cluster1::>
```

vserver services ndmp extensions show

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::>
```

vserver services ndmp log start

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 (privilege: advanced)

This parameter specifies the name of the Vserver.

`-session-id <text>` - Session Identifier (privilege: advanced)

This parameter specifies the NDMP session-id on which logging needs to be started.

`-filter <text>` - Level Filter (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
```

vserver services ndmp log stop

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 (privilege: advanced)

This parameter specifies the name of the Vserver.

`-session-id <text>` - Session Identifier (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
```

vserver services ndmp restartable-backup delete

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 (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 restartable-backup delete -vserver cluster1-
01 -context-id 0f8f5c44-d540-11e5-8c45-005056963504
    cluster1::>
```

vserver services ndmp restartable-backup show

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 restartbale 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 copy was created by DUMP engine.

[-no-acls {true|false}] - Is NO_ACLS Set? (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 correspoding 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 (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 (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? (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? (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 (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 (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 (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 (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 enviroment 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? (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? (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? (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 copy.

[-is-dp-vol {true|false}] - Is DP Volume? (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.

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.
```

vserver services web modify

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
```

vserver services web show

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 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

- vserver services web access show

vserver services web access create

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 security login show , must be included in this configuration.

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

- security login show

vserver services web access delete

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.

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 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

- vserver services web access create

vserver services web access show

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 security login show , must be included in this configuration.

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

- security login show

vserver smtape commands

vserver smtape break

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
```

vserver snapdiff-rpc-server commands

vserver snapdiff-rpc-server off

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 (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
```

vserver snapdiff-rpc-server on

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 (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
```

vserver snapdiff-rpc-server show

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] }

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

[-vserver <vserver name>] - Vserver (privilege: advanced)

If you specify this parameter, the command displays information only about the specified Vserver.

[-state {on|off}] - SnapDiff RPC Server state (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.
```

vserver vscan commands

vserver vscan disable

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.

> This command is not supported on a Vserver with Infinite Volume.

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
```

vserver vscan enable

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.

This command is not supported on a Vserver with Infinite Volume.

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
```

vserver vscan reset

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.

This command is not supported on a Vserver with Infinite Volume.

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::>
```

vserver vscan show-events

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

This command is not supported for a Vserver with Infinite 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.

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

vserver vscan show

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

This command is not supported for a Vserver with Infinite 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

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.
```

vserver vscan connection-status show-all

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

This command is not supported for a Vserver with Infinite 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.

[-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 clustered Data 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 clustered Data 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
```

vserver vscan connection-status show-connected

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

This command is not supported for a Vserver with Infinite 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.

[-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 clustered Data 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 clustered Data 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
```

vserver vscan connection-status show-not-connected

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

> This command is not supported for a Vserver with Infinite 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.

[-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
```

vserver vscan connection-status show

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

This command is not supported for a Vserver with Infinite 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.

[-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.
```

vserver vscan on-access-policy create

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.

This command is not supported for a Vserver with Infinite Volume.

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 whether access to a file is allowed if there are no external virus-scanning servers available for virus scanning. By default, it is on.

[-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...".

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.

> 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.

> 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 "?".

> 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
```

vserver vscan on-access-policy delete

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver vscan on-access-policy disable

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy enable

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy modify

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.

This command is not supported for a Vserver with Infinite Volume.

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...".

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.

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.

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 "?".

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
```

vserver vscan on-access-policy show

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

This command is not supported for a Vserver with Infinite 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

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
```

vserver vscan on-access-policy file-ext-to-exclude add

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy file-ext-to-exclude remove

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy file-ext-to-exclude show

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

This command is not supported for a Vserver with Infinite 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

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.
```

vserver vscan on-access-policy file-ext-to-include add

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy file-ext-to-include remove

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.

This command is not supported for a Vserver with Infinite Volume.

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*
```

vserver vscan on-access-policy file-ext-to-include show

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

This command is not supported for a Vserver with Infinite 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

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.
```

vserver vscan on-access-policy paths-to-exclude add

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy paths-to-exclude remove

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-access-policy paths-to-exclude show

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

> This command is not supported for a Vserver with Infinite 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

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.
```

vserver vscan on-demand-task create

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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/…".

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.

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.

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 "?".

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>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.

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
[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
```

vserver vscan on-demand-task delete

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.

> This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver vscan on-demand-task modify

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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/…".

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.

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.

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 "?".

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>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.

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
```

vserver vscan on-demand-task run

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver vscan on-demand-task schedule

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.

This command is not supported for a Vserver with Infinite Volume.

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.

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
```

vserver vscan on-demand-task show

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

This command is not supported for a Vserver with Infinite 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

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>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.

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
```

vserver vscan on-demand-task unschedule

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan on-demand-task report delete

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver vscan on-demand-task report show

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

This command is not supported for a Vserver with Infinite 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

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>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 (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 (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
```

vserver vscan scanner-pool apply-policy

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan scanner-pool create

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 clustered Data 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 vserver vscan scanner-pool apply-policy . The default applied policy is *idle* , which means the scanner pool is inactive.

This command is not supported for a Vserver with Infinite Volume.

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 clustered Data 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 clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

[-request-timeout <[<integer>h][<integer>m][<integer>s]>] - Request Service Timeout (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>h][<integer>m][<integer>s]>] - Scan Queue Timeout (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>h][<integer>m][<integer>s]>] - Session Setup Timeout (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>h][<integer>m][<integer>s]>] - Session Teardown Timeout (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 (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

- vserver vscan scanner-pool apply-policy

vserver vscan scanner-pool delete

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.

This command is not supported for a Vserver with Infinite Volume.

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.
```

vserver vscan scanner-pool modify

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 clustered Data ONTAP and external virus-scanning server, or "Vscan server". It also specifies other parameters which are used for connection management.

This command is not supported for a Vserver with Infinite Volume.

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 clustered Data 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 clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

[-request-timeout <[<integer>h][<integer>m][<integer>s]>] - Request Service Timeout (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>h][<integer>m][<integer>s]>] - Scan Queue Timeout (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>h][<integer>m][<integer>s]>] - Session Setup Timeout (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>h][<integer>m][<integer>s]>] - Session Teardown Timeout (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 (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
```

vserver vscan scanner-pool resolve-hostnames

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 addressses. 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.

This command is not supported on a Vserver with Infinite Volume.

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
```

vserver vscan scanner-pool show-active

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

This command is not supported for a Vserver with Infinite 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

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>h][<integer>m][<integer>s]>] - Request Service Timeout (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>h][<integer>m][<integer>s]>] - Scan Queue Timeout (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>h][<integer>m][<integer>s]>] - Session Setup Timeout (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>h][<integer>m][<integer>s]>] - Session Teardown Timeout (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 (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.
```

vserver vscan scanner-pool show

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

This command is not supported for a Vserver with Infinite 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

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>h][<integer>m][<integer>s]>] - Request Service Timeout (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>h][<integer>m][<integer>s]>] - Scan Queue Timeout (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>h][<integer>m][<integer>s]>] - Session Setup Timeout (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>h][<integer>m][<integer>s]>] - Session Teardown Timeout (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 (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
```

vserver vscan scanner-pool privileged-users add

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.

This command is not supported for a Vserver with Infinite Volume.

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 the privileged user or users that you want to add to the specified scanner pool.

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
```

vserver vscan scanner-pool privileged-users remove

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan scanner-pool privileged-users show

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

This command is not supported for a Vserver with Infinite 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

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.
```

vserver vscan scanner-pool servers add

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan scanner-pool servers remove

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.

This command is not supported for a Vserver with Infinite Volume.

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
```

vserver vscan scanner-pool servers show

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

This command is not supported for a Vserver with Infinite 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

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.
```

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

https://www.netapp.com/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.

https://www.netapp.com/company/legal/trademarks/

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf

Privacy policy

https://www.netapp.com/company/legal/privacy-policy/

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

- Notice for ONTAP 9.3

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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