Hewlett Packard
Enterprise

Command Reference

Published February, 2021
Version 5.3

Legal Notices

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Publication Date

Tuesday February 16, 2021 15:36:23

Document ID

ghf1602207026212

Support

All documentation and knowledge base articles are available on HPE InfoSight at https://infosight.hpe.com. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: support@nimblestorage.com

For all other general support contact information, go to https://www.hpe.com/us/en/services/nimble-storage.html.

Contents

The Nimble CLI

Use NimbleOS administrative commands from the command line interface (CLI) to configure and monitor an HPE Nimble Storage array.

Use a secure shell (SSH) utility to log in to the management IP of an HPE Nimble Storage array to run administrative commands. Your role determines your permission level, which in turn determines the command options you can run.

After you log in, you see a security message for array usage. You do not have to acknowledge the security message.

Enter ? to display a list of administrative commands. Then use any or all of the following options to get more information about the commands.

- To display available options for a specific command, enter command_name --help.

  The output lists each option and suboption, and provides brief descriptions of each.

- To display a man page that introduces all the commands, type man intro.

  The output lists each command, and provides brief descriptions of each. Press the space bar to page through the list. Press q to quit when you get to the end of the introduction.

- To display detailed information for a specific command, enter man *command_name*.

  The output lists the command name and brief description, synopsis of each option and suboption, and detailed descriptions of each option and suboption, including valid values. Press the space bar to page through the list. Press q to quit when you get to the end of the man page.

CLI Help

The ? command lists the NimbleOS administrative commands or all commands.

> **Note:** The ? and help commands are similar. The major difference is that help command has a man page; the ? command does not.

Synopsis

?

? --help

? --all

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--all	N/A	Show all available commands, not just the Nimble-specific administrative commands.	Guest

Examples

This example lists the administrative commands used to manage a Nimble array.

```
Nimble OS $ ?
?                   fc                  pool                subnet
alert               group               prottmpl            timezone
array               halt                reboot              useradmin
auditlog            help                route               usersession
cert                initiatorgrp        setup               version
chapuser            ip                  shelf               vmwplugin
ctrlr               migration           snap                vol
date                netconfig           snapcoll            volcoll
disk                nic                 software
encryptkey          partner             sshkey
failover            perfpolicy          stats
```

This example shows partial output from ? --all. The output is truncated to save space.

```
Nimble OS $ ? --all
?                   expr                nohup               subnet
[                   failover            nslookup            sum
[[                  false               od                  sync
alert               fc                  partner             tac
arp                 fgrep               patch               tail
array               find                perfpolicy          tee
...
...
encryptkey          mv                  stats               zcat
env                 netconfig           strings
expand              nic                 stty
```

Nimble Administrative Commands

A wide variety of administrative commands are provided to configure, manage, and monitor HPE Nimble Storage arrays.

Use either the NimbleOS graphical user interface (GUI) or the command line interface (CLI) to perform administrative operations on an HPE Nimble Storage array or group of arrays.

These interfaces provide comparable management capabilities. The GUI has tool tips, embedded user assistance, and help to guide you through administrative operations interactively. The CLI is useful to script frequently performed operations.

alarm

The alarm command lets you view and manage alarms and notifications. By default, all alarms are displayed when you use the alarm command with the --list option. However, you can filter for severity levels, categories, and time spans. For example, you might want to generally limit alarms to only recent ones, but regularly view a longer-term list of alarms to look for patterns of error. Filtering lets you manage only the alarms you are interested in. You can also add and modify reminders or notifications to alarms such as changing the frequency of alarm notifications.

Synopsis

alarm --help

alarm --list
[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--category *name*]
[--severity *level*]

alarm --info *id*

alarm --acknowledge *id*
[--remind_every *period*]
[--remind_every_unit {minutes|hours|days}]

alarm --unacknowledge *id*

alarm --edit *id*
[--remind_every_period]
[--remind_every_unit {*minutes|hour|days*}]

alarm --delete *id*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--list	N/A	List all alarms and their IDs. For example alarm --list shows all info, while alarm --list --from 08-14,11:26:00 lists only alarms that appear after August 14th at 11:26 am. You can specify time with no date, but if you specify a date you must include a time. The --list option takes the following sub-options.	Operator
	[--from *yyyy-]mm-dd,]hh:mm[:ss]*]	List alarms starting from a specified time. Often used with the --to option to frame an alarm list. Use only the time, or the date and time, for example, alarm --list --from 2015-08-25,12:10:00 or alarm --list --from 12:10:00.	
	[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]	List alarms until specified time. Often used with the --from option. If not specified,--list continues with no end time. Use only the time, or the date and time, for example, alarm --list --to 2017-08-25,12:10:00 or alarm --list --from 05:15 --to 12:10:00.	
	[--category *name*]	Restrict the list of alarms to only include the specified category. Valid categories are hardware, service, replication, volume, update, config and security.	
	[--severity *level*]	Restrict the list of alarms to only include the specified severity. Valid severities are warning and critical. The option will only return alarms of the severity specified.	
--info	*id*	Show detailed information about a specified alarm. For example, alarm --info 668 displays information about a alarm 668. To determine an alarm ID, enter alarm --list.	Guest

Option	Arguments and Suboptions	Description	Role
--acknowledge	*id*	Acknowledge an alarm with the current username. For example, alarm --acknowledge 668 makes the current user the one that acknowledges alarm 668, and sets the acknowledged date as the current date. To determine the alarm ID, enter alarm --list. The --acknowledge option takes the following sub-options.	Guest
	[--remind_every *period*]	Frequency of alarm notification. The valid range is 0-60. If this option is unspecified, it defaults to *1* and --remind_every_unit is set to *day*, so an alarm notification is sent every hour. If the --remind_every option is set to 0, a reminder email will not be sent.	
	[--remind_every_unit {minutes|hours|days}	Units for reminder frequency in minutes, hours or days. If this option is unspecified, the default unit of *day* is used.	
--unacknowledge	*id*	When you unacknowledge an alarm, it clears the acknowledged user and date fields. To determine the alarm ID, enter alarm --list.	Operator
--edit	*id*	Edit the notification frequency for a specified alarm. To determine the alarm ID, enter alarm --list. The --edit option takes the following sub-options.	Operator
	[--remind_every *period*]	The frequency of alarm notifications. The valid range is 0-60. If this option is unspecified, it defaults to *1* and --remind_every_unit is set to *day*, so that an alarm notification is sent every hour. If the --remind_every option is set to 0, a reminder email will not be sent.	
	[--remind_every_unit {minutes|hours|days}	Units for reminder frequency in minutes, hours or days. If this option is unspecified, the default unit of *day* is used.	
--delete	*id*	Delete a specified alarm. For example, alarm --delete 1 deletes the current alarm *1* if it exists. Be careful using this command because an alarm indicates severe conditions on the system and should not be ignored. To determine the alarm ID, use alarm --list.	Power user

Example

The following example lists all alarms in the group where you run the command.

```
Nimble OS $ alarm --list
-----+----------+--------------------+---------+--------+------------------------------
ID     Severity       Time               Status    Array    Detail
-----+----------+--------------------+---------+--------+------------------------------
5      CRITICAL  Oct 20 2016 16:49:30    Open      -       Volume test string
 from alert
gen space usage is at 0% and approaching quota of 0%. It will be taken offline
 if it exceeds
the quota
```

The following example lists all the alarms on the array (there is only one) and then acknowledges that alarm setting the notifications of the alarm to remind the user every day.

```
Nimble OS $ alarm --list
-----+----------+--------------------+---------+---------+------------------------------
ID     Severity       Time               Status    Array     Detail
-----+----------+--------------------+---------+---------+------------------------------
5      CRITICAL  Oct 20 2016 16:49:30     Open     -        Volume test string
 from alert
gen space usage is at 0% and approaching quota of 0%. It will be taken offline
 if it exceeds
the quota

Nimble OS $ alarm --acknowledge 5 --remind_every 1 --remind_every_unit days
```

The following example lists detailed information about the alarm with ID 5.

```
Nimble OS $ alarm --info 5
ID: 5
Severity: critical
Category: volume
Onset time: Oct 20 2016 16:49:30
Array: -
Object type: volume
Object name: test
Description: Volume test string from alert gen space usage is at 0% and
approaching quota of 0%.
It will be taken offline if it exceeds the quota
Status: Acknowledged
Acknowledged by: admin
Acknowledged time: Oct 26 2016 15:59:55
Repeat period: Every 1 days
Next reminder time: Oct 27 2016 16:00:48
```

alert

The alert command options manage array alerts, which include certain events. You can also use the command options to display information about alerts and events. Alerts and events can use both email and SMTP notification methods.

> **Note:** The terms *event* and *alert* are sometimes used interchangeably. While an event can have different severity levels, the term *event* is often used to refer to a non-critical issue, such as a change to your environment. The term *alert* refers to real-time notifications about issues, such as warning events, that can adversely affect your system. You can use alerts and events as a diagnostic and preventative tool. For more information about alerts and events, see either the *GUI Administration Guide* or the *CLI Administration Guide*.

By default, all alerts and events are displayed when you use the alert --list command. You can filter the display based on severity levels, categories, and time spans so that you display only the alerts that you are interested in. For example, you can filter for the most recent alerts, and periodically list all of them to look for potential errors.

You can also send a test alert to verify that the system is configured properly to send alerts.

Synopsis

alert --help

alert --list
[--array *array_name*]
[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--severity *{info|warning|critical}*]
[--category *{hardware|service|replication|volume|update|config}*]

alert --info *id*

alert --test
[--severity *{info|notice|warning|critical}*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	[--array *array_name*]	List alerts and events for the specified array.	Guest
	[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]	List alerts that start from the specified time. Use only the time, or use the date and time. For example, alert --list --from 08-14,12:10:00 or alert --list --from 12:10:00. Use this option with or without the --to option. If used alone, --to is assumed to be now.	
	[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]	List alerts until the specified time. Use only the time, or use the date and time. For example, alert --list --to 08-14,12:10:00 or alert --list --to 12:10:00. Use this option with or without the --from option.	
	[--severity *{info|notice|warning|critical}*]	List alerts and events of the specified severity level. Alert levels are not cumulative. For example, selecting critical does not show the info or warning severity levels.	
	[--category *{hardware|service|replication|volume|update|config|security}*]	List alerts and events of the specified category.	

Option	Arguments and Suboptions	Description	Role
--info	*id*	List detailed information about the specified alert ID. To determine valid alert IDs, use alert --list.	Guest
--test	--severity *{info|notice|warning|critical}*]	Generate a test alert with the specified severity level.	Power User

Examples

This example lists all alert messages on the array where you run the command. The output is truncated to save space.

```
Nimble OS $ alert --list
----------+--------+--------------------+------+-----------+----------+---------
ID         Severity Time                 Type   Category    Array      Detail
----------+--------+--------------------+------+-----------+----------+---------

         8 INFO     Jun 27 2014 11:13:40   2101 service     AA-102081  System
 services started...
         9 INFO     Jun 27 2014 11:14:02  12714 hardware    AA-102081  A new
shelf is detected...
...
```

This example lists information about the alert with the ID 23.

```
Nimble OS $ alert --info 23
ID: 23
Active: yes
Priority: INFO
Category: hardware
Type: 12201
Time: Feb 24 2014 17:10:00
Cleared: N/A
Array: AC-102266
Target type: NIC
Target: 192.168.1.3
Detail: IP interface 192.168.1.3 up on controller A NIC port eth1
```

This example sends an alert test email with the warning severity level.

```
Nimble OS $ alert --test --severity warning


From: Admin
Sent: Wednesday, August 14, 2013 12:12 PM
To: eng-sol
Subject: Nimble Alert on sol / mars (AC-109828) - WARNING: Test alert
Time: Wed Aug 14 15:11:40 2013

Type: 5005
Id: 31645
Message: Test message at warn level

Group Name: Admin
Group ID: 2070866469188534807
Array name: mars
Serial: AC-109828
Version: 2.0.3.8-46920-opt
```

```
Arrays in the group:
----------------+-----------------+-----------+-----------------+-----------
Name            Serial            Model       Version           Status
----------------+-----------------+-----------+-----------------+-----------
moon            AC-103234         CS460G-X2   2.0.3.8-46920-opt reachable
saturn          AC-103236         CS460G-X2   2.0.3.8-46920-opt reachable
jupiter         AC-103242         CS460G-X2   2.0.3.8-46920-opt reachable
mercury         AC-103235         CS460G-X2   2.0.3.8-46920-opt reachable
```

array

The array command options manage a Nimble array.

> **Note:** Minimize the amount of time between running array --resetup and the setup command. On a Fibre Channel array, there is no I/O between running array --resetup and setup. On an iSCSI array, existing connections may be unaffected, but new connections are refused.

Synopsis

array --help

array --list

array --discover

array --info *{name|serial_number}*

array --edit *name*
[--name *new_name*]

array --resetup
[--non_interactive]

array --add *serial_number*
[--name *name*]
[--subnet_label *subnet_label*]
[--data_ipaddr *ipaddr*]
[--nic *nic*]
[--tagged *tag*]
[--support_ipaddr *ipaddr*]
[--pool *pool*]
[--create_pool]
[--pool_description *pool_description*]
[--allow_lower_limits]

array --remove *name*
[--force]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--list	N/A	List all discovered and group member arrays.	Guest
--discover	N/A	Discover non-member arrays on the same management subnet.	Power User
--info	{*name* | *serial_number*}	Provide detailed information about the specified group-member array.	Guest
--edit	[--name *new_name*]	Modify the name of the array. The name must conform to iSCSI naming conventions. It can be any combination of up to 63 alphanumeric characters and hyphen, but cannot begin with a hyphen, or contain spaces or other punctuation.	Power User
--resetup	N/A	Set a standalone array in a state where you can run the setup command again. This reenables the zeroconf process and lets you reset the array password, IP address, and most other configuration settings. No data will be lost. After running this command, any computer on the same subnet as the array can be used to reset the array password and then control the array. Use this command with caution. Minimize the time between resetting the array with this command and running setup again to limit I/O disruption. On a Fibre Channel array, there is no I/O between running array --resetup and setup. On an iSCSI array, existing connections may be unaffected, but new connections are refused.	Administrator
	[--non_interactive]	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly setting up the array.	

Option	Arguments and Suboptions	Description	Role
--add	*serial_number*	Serial number of a discovered array to add to a group.	Power User
	[--name *name*]	Name of a array to add to a group. The group must exist.	
	[--subnet_label *subnet_label*]	Subnet label on the NIC. Repeat this option to specify more than one subnet label. The first subnet label is assigned to the first NIC, such as eth1, the second subnet label is assigned to the second NIC, such as eth2, and so on. To skip a subnet label, enter a blank argument.	
	[--data_ipaddr *ipaddr*]	The network IP address(es) that will manage data traffic. Repeat this option to set multiple IP addresses for multiple NICs. Then the first IP address is assigned to the first NIC, such as eth1, and the second to the second NIC, such as eth2, and so on. To skip a NIC, type a blank argument.	
	[--nic *nic*]	NIC and subnet setting.	
	[--tagged *tag*]	Specify tagged assignment. Repeat this option to specify more than one assignment. The first tag is assigned to the first NIC, such as eth1, the second tag is assigned to the second NIC, such as eth2, and so on. To skip an assignment, enter a blank argument.	
	[--support_ipaddr *ipaddr*]	Include this argument twice. Then the first IP address entered is assigned to Controller A and the second to Controller B. You must assign static IP addresses to allow array access if the management IP address is not accessible.	
	[--pool *pool_name*]	Pool to assign the array to. If you specify this suboption, then the pool must exist. If you do not specify this suboption or the --create_pool suboption, then the array is assigned to the default pool.	
	[--create_pool]	Create a new storage pool and automatically assign the array to it. If you do not specify this suboption along with the --pool suboption, then the array is assigned to the default pool.	
	[--pool_description *pool_description*]	If you create a new pool for an array, type a description for the pool. If there are spaces in the description, enclose the text in quotation marks.	
	[--allow_lower_limits]		

Option	Arguments and Suboptions	Description	Role
		If you add an array to an existing pool and the pool's limits are higher than the array's limits, the resulting pool will have the lower limits of the array being added, which could also change the group limits to lower values.	
--remove	*name*	Name of the array to remove from a group. Use the pool command to remove the array from any pools.	Power User
	[--force]	Forcibly remove the specified array from a group even if it is unreachable. This may lead to loss of data if a volume is currently striped across disks on the array.	

Examples

This example displays a list of member arrays.

```
Nimble OS $ array --list
-----------------+---------------+-----------+-----------------+-----------
Name              Serial          Model       Version           Status
-----------------+---------------+-----------+-----------------+-----------
c12-array3        AC-100159       CS210       2.0.4.1-44556-opt reachable
c12-array6        AC-100209       CS220       2.0.4.1-44556-opt reachable
```

This example discovers non-member arrays.

```
Nimble OS $ array --discover
-------------+----------+-----------------+--------------------------------
Serial Number Model      Version           Link-Local IP Addresses
-------------+----------+-----------------+--------------------------------
AC-100159     CS210      2.0.3.8-44640-opt 169.254.8.98, 169.254.8.93...
```

This example shows information for the array named *c20-array1*.

```
Nimble OS $ array --info c20-array1
Model: CS220
...
Array name: c20-array1
...
Member of pool: default
Status: reachable
```

This example shows information for the array with the serial number *AC-102566*.

```
Nimble OS $ array --info AC-102566
Model: CS220
Serial: AC-102566
...
Array name: c20-array1
...
Member of pool: default
Status: reachable
```

This example renames the *MKTG* array to *Marketing*.

```
Nimble OS $ array --edit MKTG --name Marketing
```

This example resets a standalone array so the setup command can be run.

> **Note:** Use this command option with caution.

```
Nimble OS $ array --resetup
WARNING: This operation will reset the array configuration. Minimize the time
 before
running setup again to limit I/O disruption. Existing iSCSI connections may
be
unaffected, but new iSCSI connections will be refused.
Type yes to continue, no to cancel: yes
Successfully initiated array reset.
Minimize the time before running setup again to limit I/O disruption. Existing
 iSCSI
connections may be unaffected, but new iSCSI connections will be refused.
WARNING: While the volume data and some of the configuration data are preserved,
 a
future setup of the array could change the group name and the network
configuration.
This will invalidate the replication partner and the host configuration, which
 may
disrupt replication and iSCSI connections to this array.
```

This example adds the serial number of an uninitialized array and assigns a new name to the array during this process.

Although only one set of name, subnet_label, data_ipaddr, and pool options is required, you can use multiple sets of these options to specify values for more than one NIC on an array, such as eth1, eth2, tg1, tg2.

```
Nimble OS $ array --add AC-100159 --name  c12-array3
       \
      --data_ipaddr '' --subnet_label management      \
      --data_ipaddr '' --subnet_label management      \
      --data_ipaddr 198.51.100.61 --subnet_label data \
      --data_ipaddr 198.51.100.62 --subnet_label data \
      --support_ipaddr 10.19.0.57 --support_ipaddr 10.19.0.58
```

This example removes an existing member array named *greydient* from a group.

> **Note:** The array to remove must not be assigned to a pool. If needed, run pool --unassign to remove the array from a pool before you run this command.

```
Nimble OS $ array --remove greydient
```

auditlog

The auditlog command options manage array audit log entries.

The NimbleOS audit log keeps records of all non-read operations performed on the array, and which user performed the operation. You can search the audit log by activity and object type/name. You can also filter the audit log by time range, username, userid, activity category, and access type. Administrators can view the audit log in a summary table with faceted browsing by time, activity category, and across access type.

Synopsis

auditlog --help

auditlog --list
[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--username *username*]
[--userid *userid*]
[--category *{provision|protection|dataaccess| useraccess|sysconfig|swupdate}*]
[--access_type *{gui|cli|api}*]

auditlog --info *id*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all audit log records.	Administrator
	[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]	List audit logs that start from the specified time. Use only the time, or use the date and time. For example, auditlog --list --from 08-14,12:10:00 or auditlog --list --from 12:10:00. Use this option with or without the --to option. If used alone, --to is assumed to be now.	
	[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]	List audit logs until the specified time. Use only the time, or use the date and time. For example, auditlog --list --to 08-14,12:10:00 or auditlog --list --to 12:10:00. Use this option with or without the --from option.	
	[--username *{username}*]	List audit logs by username.	
	[--userid *{userid}*]	List audit logs by userid.	
	[--category *{provision|protection|dataaccess| useraccess|sysconfig|swupdate}*]	List audit logs of the specified category.	
	[--access_type *{gui|cli|api}*]	List audit logs using the specified access type.	
--info	*id*	List detailed information about the specified audit log ID. To determine valid audit log IDs, use auditlog --list.	Administrator

Examples

This example lists all alert messages on the array where you run the command. The output is truncated to save space.

```
Nimble OS $ auditlog --list
--------+--------------------+--------------+---------+----------------------
```

```
ID       Time                  User            Status    Activity
--------+---------------------+---------------+---------+-----------------------
1        Mar 25 2015 09:10:52 admin           Succeeded Login attempt

2        Mar 25 2015 09:10:53 admin          Succeeded Complete setup on  ...
3        Mar 25 2015 09:45:52 admin          Succeeded User admin session ...
4        Mar 26 2015 13:22:41                 Succeeded Logout user

5        Mar 26 2015 13:43:21                 Succeeded Logout user

6        Mar 27 2015 13:17:22                 Succeeded Logout user

7        Mar 27 2015 14:24:22                 Succeeded Logout user
```

This example lists all alert messages on the array where you run the command, filtered by username. The output is truncated to save space.

```
Nimble OS $ auditlog --list--username pAdmin
--------+---------------------+---------------+---------+-----------------------
ID       Time                  User            Status    Activity
--------+---------------------+---------------+---------+-----------------------

 57      Apr 22 2015 11:45:20 pAdmin          Succeeded Login attempt

 78      Apr 22 2015 13:16:40 pAdmin          Succeeded Logout user pAdmin

105      Apr 22 2015 14:44:51 pAdmin          Succeeded Login attempt

108      Apr 22 2015 14:45:47 pAdmin          Succeeded Logout user pAdmin

109      Apr 22 2015 14:46:39 pAdmin          Succeeded Login attempt

110      Apr 22 2015 14:47:19 pAdmin          Succeeded Logout user pAdmin

403      Apr 23 2015 10:37:16 pAdmin          Succeeded Login attempt

410      Apr 23 2015 11:05:15 pAdmin          Succeeded Logout user pAdmin

455      Apr 23 2015 14:11:37 pAdmin          Succeeded Login attempt

456      Apr 23 2015 14:12:25 pAdmin          Succeeded Logout user pAdmin

707      Apr 24 2015 07:23:12 pAdmin          Succeeded Login attempt
```

This example lists information about the audit log with the ID 23.

```
Nimble OS $ auditlog --info 23
ID: 23
Type: 8015
Object name: vol2
Object type: Volume
Scope: -
Time: Mar 25 2015 09:28:06
Status: Succeeded
Error code: SM_http_ok
User ID: 3
Username: pAdmin
User full name: Administrator User
Source IP: 10.20.21.168
Access Type: GUI
```

```
Application name: GUI
Category: Data Protection
Activity type: other
Activity: Associate volume vol2 with volume collection volcoll2
```

cert

The cert command is used to manage certificates used by a number of components that support SSL/TLS. These include the web GUI interface, the VASA provider, the CIM server, and Group Management REST APIs. It is possible to select different certificate chains for each of these. In addition, a list of trusted certificates can be managed for use by internal clients, for example to provide verification of an external key manager that presents its own self-signed certificate on connection.

To establish a secure connection with a website or other server, the array presents a certificate to authenticate its identity. Certificates are an important component of the Secure Sockets Layer (SSL), because they prevent others from impersonating a secure website or other server.

The default values used for the *subject*, *dnslist*, and *iplist* are taken from the existing group certificate, if no corresponding options are provided.

There are three certificate chains that are managed using the cert command: *array*, *group*, and *custom*.

Array certificate chain:

The array certificate is generated automatically before group setup.

Group certificate chain:

The group certificate is generated after the group leader has been configured.

Custom certificate chain:

A custom certificate may also be installed and used. It can be either a self-signed (*custom*), or a Certificate Signing Request (*custom-csr)* that can be exported to be signed by a local Certificate Authority (CA). The CA certificate and the resulting signed custom certificate may then be imported for use.

Synopsis

cert --help

cert --list

cert --info *cert-name*

cert --gen {*array|group|custom|custom-csr*}
[--subject *subject_name*]
[--dnslist *dns_list*]
[--iplist *ip_list*]
[--num_days *days*]
[--check]
[--force]

cert --import *cert-name*
[--pks12]
[--password *unlock_password*]
[--trusted]
[--from_host *host:port*]
[--force]

cert --delete {*cert-name*}

cert --use {*array|group|custom*}
[--https]
[--apis]

cert --verify *cert-name*
[--checkend *days*]
[--oscp]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Shows usage information for the command.	Administrator
--list	N/A	List each of the certificates installed, and which are being used for HTTPS and API access.	Administrator
--info	{*cert-name*}	Print output that shows all of the details, as well as the PEM encoding, of the given certificate or CSR.	Administrator

Option	Arguments and Suboptions	Description	Role
--gen	{*array*|*group*| *custom*|*custom-csr*}	Initiates the creation of either a self-signed certificate, or a Certificate Signing Request (CSR). The argument specifies whether to create a certificate of the type specified or to create a CSR for a custom certificate.The *custom* option generates a self-signed certificate with the parameters specified, while the *custom-csr* option generates a CSR with the parameters specified. A custom certificate and be either self-signed or signed by an external Certificate Authority (CA). The *array* and *group* certificates are generated automatically for use by HTTPS and API services.	Administrator
	[--subject *subject-name*]	Subject or common name for the certificate. The default value for the subject name will be the group name. If the value is "-" or the option is not present, then an attempt is made to extract the subject name from the existing group certificate, if there is one. The value is treated as a full subject name if it begins with "/", in the format that OpenSSL uses. Otherwise, it is taken as the Common Name value.	
	[--dnslist *dens-list*]		

Option	Arguments and Suboptions	Description	Role
		A comma separated list of DNS names to add to the subject alternate name field of the host certificate. If the value is "-" or the option is not specified, an attempt is made to extract the DNS names from the existing group certificate, if there is one.	
	[--iplist *ip-list*]	A comma separated list of IP addressess to add to the subject alternate name field of the host certificate. If the value is "-" or the option is not specified, an attempt is made to extract the DNS names from the existing group certificate, if there is one. To specify an empty IP list, use "--ip-list ""	
	[--num_days *days*]	Number of days for which the certificate will be valid. The default value of the number of valid days is 3650.	
	[--check]	Check the current subject, DNS list, and IP list of the current group certificate against the provided inputs. A new certificate chain is only generated if the checked values do not match the group certificate values.	
	[--force]	Force overwriting of an existing certificate chain. Otherwise, the existing certificate is left as is.	

Option	Arguments and Suboptions	Description	Role
--import	*cert-name*	Import a custom CA (*custom-ca*), signed custom certificate (*custom*), or a trusted certificate to be used by internal clients. The custom-ca certificate must be imported before the signed custom certificate. After the custom certificate is imported and verified, it will be installed for use. The imported certificate should be valid for use as both SSL server and client, since it may be used as an outbound client certificate. A combined bundle of key, CA certificate, and signed certificate may also be imported using the PKCS12 format. This format is most useful for installing a wildcard certificate.	Administrator
	[--pkcs12]	Import a PKCS12 file that contains a private key, CA certificate(s), and signed host certificate. Use the custom command argument.	Administrator
	[--password *unlock-password*]	The password to unlock a PKCS12 input file.	Administrator
	[--trusted]	Import a certificate to add to the list of certificates trusted by internal SSL clients. The certificate can either be read from standard input, or retrieved from the specified host. When this flag is set, the names custom-ca, custom, group, and array are reserved.	Administrator

Option	Arguments and Suboptions	Description	Role
	[--from_host *host:port*]	When used with the --trusted option, attempt an SSL connection to the specified host and port, and import the presented certificate to the trusted certificate list.	Administrator
	[--force]	Force overwriting of an existing certificate. Otherwise, the existing certificate is left as is.	
--delete	{*cert-name*}	Delete a certificate, along with any associated private key and optional certificates. Deleting the array or array-ca certificate is not allowed.	Administrator
--use	{*array|group|custom*} [--https] [--apis]	Select one of the array, group, or custom certificate chains to use for HTTPS access or API use. Use the named certificate for HTTPS service. Use the named certificate chain for access to Nimble APIs (REST and internal Nimble APIs).	Administrator

Option	Arguments and Suboptions	Description	Role
--verify	*cert-name*	Verify the signature, expiration time, and (optionally) the OCSP status of the named certificate. When used with no suboptions, verify the signature and whether the certificate has expired.	Administrator
	[--checkend *days*]	Check if the certificate expires withing the specified number of days. If not present, the default value is 0.	Administrator
	[--oscp]	If the certificate as an OSCP reference included, contact the OSCP service and determine the status of the certificate.	Administrator

Examples

To create and install a custom certificate on an array, you must perform the following steps from an ssh session:

Nimble OS $ cert --gen custom-csr --subject '/C=US/ST=CA/L=San Jose/O=Nimble Storage/OU=Engineering/CN=AF106656' --dnslist group-kent-vma.nimblestorage.com,kent-vma.nimblestorage.com

```
-----BEGIN CERTIFICATE REQUEST-----
MIICuTCCAaECAQAwWDELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMREwDwYDVQQH
DAhTYW4gSm9zZTEXMBUGA1UECgwOTmltYmxlIFN0b3JhZ2UxEDAOBgNVBAMMB2xu
LWtlbnQwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQD1p+rdq/GvuALc
qqam6AIoKfC8X4WFN+xOSmco7cQliZnmYNyX1Rzrw0RyIYwfBOeSff+47pw6Hv6S
TY54lMMf7/5vMmEoR56DaolkI9RgmMnQM8Ay6YnRjWZtE7TFRiId+r7cYlPlREWd
0Mq25psKqEaNl8QzNVFhxOdPCe//cP6Z91R2GanR+cafAmSZ8ELfkFvdlWnU5Vdv
cCHH73IkHtZbgWVy9yJJ2ZZT9Lv31rPjx2rP+h305VrWWa7xWyTI0FD86y5n41Zn
Yx8ypE8E0dcSENlNX9evRzGyPfhIg11JaHTkad/REpAwXSDilEcl4/0jRfKBaq/M
YD6U1Z8tAgMBAAGgHDAaBgkqhkiG9w0BCQ4xDTALMAkGA1UdEQQCMAAwDQYJKoZI
hvcNAQEFBQADggEBABcwx0s4dvaBbVv2NYNs/tYq/pwynGcF19+MIVtrCxoYQ3h3
MObHTu0k0+LgWNZeJKDfD2AKHGnQMjP1lWqHqbZLwxbioFotxIYd/XwIMSlshBLS
WegYzrCbfDTeQMLV7XozIfcGEUZzQF+hSY7SPJGLRZTHCAI/hD0LXthl6fUB2ynZ
QsjA3q38+OV/H37vkXA0jY5oxFcPwPGGh6jLVMBspY2OHnxGiUIEganpxk3XtOAh
bYdvo88B0WRV7chDfL2/VzythDL8vOPxSsY8leCdXOq8dnUb2t5Yo8cBqc2J58ON
HiDL7gI4YR/EafXDuHJwGnqMqGGlddktAQkTWWI=
-----END CERTIFICATE REQUEST-----
```

You must then cut-and-paste the certificate request output into a signing request to the CA.

Then cut-and-paste the CA certificate and the output from the signing into the command inputs below:

```
Nimble OS $ cert --import custom-ca
Nimble OS $ cert --import custom
```

```
-----BEGIN CERTIFICATE-----
MIIEADCCAuigAwIBAgIJALkB09sKkc1IMA0GCSqGSIb3DQEBBQUAMF0xCzAJBgNV
BAYTAlVTMQswCQYDVQQIDAJDQTERMA8GA1UEBwwIU2FuIEpvc2UxFzAVBgNVBAoM
Dk5pbWJsZSBTdG9yYWdlMRUwEwYDVQQDDAxrZW50YXJyYXktQ0EwHhcNMTUwNjIy
MjIxNDMxWhcNMjUwNjE5MjIxNDMxWjBdMQswCQYDVQQGEwJVUzELMAkGA1UECAwC
Q0ExETAPBgNVBAcMCFNhbiBKb3NlMRcwFQYDVQQKDA5OaW1ibGUgU3RvcmFnZTEV
MBMGA1UEAwwMa2VudGFycmF5LUNBMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIB
CgKCAQEAmLA18pIh44mkgGPll9jpey+IXso+mwJAckrfgxh8Z7+54L9K5/sGy4Fe
nN8+DsPc7qtNl1lnQ0nptkDjmQph3sV34NB1nOg+cI+pSMRC1kZk2ZAfW0xWeJJI
/JBTQFWOzbe2PgmJNMN0gMI8eeco9PmR3n9+WEGsos6SsUsW/WHSNo41ujuZV0wm
LmLP/mv1WuoUWpDZwyzk0uovSihkcTbCihm7yHYpKMxvYMy88/Hea7rBbQaid9z/
xceEZ6cnEotZcEyJP+u0Z7FMiPfzcPY77XV1IsiVB1eVohOP1GURlD2J9v1Xli1U
1H98R4vGeynI7rjxo7dhsZu1K9xVIwIDAQABo4HCMIG/MB0GA1UdDgQWBBShBlOm
mIP2S84mJ3wTMoWJgEYumzCBjwYDVR0jBIGHMIGEgBShBlOmmIP2S84mJ3wTMoWJ
gEYum6FhpF8wXTELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMREwDwYDVQQHDAhT
YW4gSm9zZTEXMBUGA1UECgwOTmltYmxlIFN0b3JhZ2UxFTATBgNVBAMMDGtlbnRh
cnJheS1DQYIJALkB09sKkc1IMAwGA1UdEwQFMAMBAf8wDQYJKoZIhvcNAQEFBQAD
ggEBAGTzXWDcMI/OnP/KerhpGPfns0VUMzvvYaKn5Db5ADEPL5vafawhsOyTMgnu
7hwiL11Zy6qJ+e0g85Zm0v0mg4wKyM3QKvbf86lu8YKvEoYoCWcOcsRDt2ijrbx7
IkRyXOypI0nIN76vNQpo1NtZ+e3XvzAShrFdZNtsi8ILZqEOepr/PBNJOv1hAszv
CpcZ6uBMxodez2lHT7ceb5ZxYJZnOoNkLTMKp9MV7Qt0v8xDKWMy9ED9PuJ79f+v
pzmJPTcv2pFHrMoPXqZFxXVTXk5GwjORGlbqvAdsGexZ3vqbmfgQvVuDdw9Zd60v qkg4gT6MJsnfdgsOgpRg5/4CRao=

-----END CERTIFICATE-----
```

```
-----BEGIN CERTIFICATE-----
MIIDhTCCAm2gAwIBAgIJAId2ciTqy7mCMA0GCSqGSIb3DQEBBQUAMF0xCzAJBgNV
BAYTAlVTMQswCQYDVQQIDAJDQTERMA8GA1UEBwwIU2FuIEpvc2UxFzAVBgNVBAoM
Dk5pbWJsZSBTdG9yYWdlMRUwEwYDVQQDDAxrZW50YXJyYXktQ0EwHhcNMTUwNjIy
MjIxNDMyWhcNMjUwNjE5MjIxNDMyWjBYMQswCQYDVQQGEwJVUzELMAkGA1UECAwC
Q0ExETAPBgNVBAcMCFNhbiBKb3NlMRcwFQYDVQQKDA5OaW1ibGUgU3RvcmFnZTEQ
MA4GA1UEAwwHbG4ta2VudDCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEB
APWn6t2r8a+4AtyqpqboAigp8LxfhYU37E5KZyjtxCWJmeZg3JfVHOvDRHIhjB8E
55J9/7junDoe/pJNjniUwx/v/m8yYShHnoNqiWQj1GCYydAzwDLpidGNZm0TtMVG
Ih36vtxiU+VERZ3QyrbmmwqoRo2XxDM1UWHE508J7/9w/pn3VHYZqdH5xp8CZJnw
Qt+QW92VadTlV29wIcfvciQe1luBZXL3IknZllP0u/fWs+PHas/6HfTlWtZZrvFb
JMjQUPzrLmfjVmdjHzKkTwTR1xIQ2U1f169HMbI9+EiDXUlodORp39ESkDBdIOKU
RyXj/SNF8oFqr8xgPpTVny0CAwEAAaNNMEswCQYDVR0TBAIwADAdBgNVHQ4EFgQU
dudH4pEro5KcsujvyGr728sKSyAwHwYDVR0jBBgwFoAUoQZTppiD9kvOJid8EzKF
iYBGLpswDQYJKoZIhvcNAQEFBQADggEBAJcsLauYn1AmaApF4uag6kljGls2K2qm
DNBQy2w2IWM9ffxSR+j0DhVgbgHQrMkTwCy9pgFAXS9MXCVF3BfLW1J9UdZXmegt
AsITa9xDthJqMwvweT7s3Gzwlc8zRXZOCXe79+YF3sZhWlFV5E1ex1iHmIwA7x9T
Twe4timzAsyjKQFOSm8ggQxbizoTDXpehiDgSDxzqL+8wggi13sC4ibAgGDzGAvC
u4EwtNlluIKBGnDHaGyfxv8ZT0jkHEB/KTFZAq8a0TVRZHi9TpnGszNcA9aV8ZKC
VHZPdBqjcDJBjFiLFZa9palRHIl3Q3laLVU+lEQZhD0M/x0TZwvPlUY=
-----END CERTIFICATE-----
```

The following example is used to change APIs to use the group certificate.

Display the current status.

```
sjc-array289-A:/# cert --list
Name        Subject
====================
array:      /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=AF-106720
group:      /C=US/ST=CA/L=San Jose/O=Nimble
Storage/CN=sjc-array289.WIN2K8R2AD.NET
custom:     /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=sjc-array289-A
custom-csr: /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=sjc-array289-A

Use         Name
====================
HTTPS:      custom
APIs:       custom
```

Run the command to change APIs to use the group certificate.

```
sjc-array289-A:/# cert --use group --apis
INFO: group certificate chain installed for API use
```

Run the command to display the updated status.

```
sjc-array289-A:/# cert --list
Name        Subject
====================
array:      /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=AF-106720
group:      /C=US/ST=CA/L=San Jose/O=Nimble
Storage/CN=sjc-array289.WIN2K8R2AD.NET
custom:     /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=sjc-array289-A
custom-csr: /C=US/ST=CA/L=San Jose/O=Nimble Storage/CN=sjc-array289-A

Use         Name
====================
HTTPS:      custom
APIs:       group
```

chapuser

Synopsis

The chapuser command options manage Challenge-Response Handshake Authentication Protocol (CHAP) user accounts. CHAP users are one method of access control for iSCSI initiators. Each CHAP user has a CHAP password, sometimes called a CHAP secret.

CHAP passwords must be the same on the array and the iSCSI initiator before the array can authenticate the iSCSI initiator and permit access. The CHAP user information must exist on both the array and the iSCSI initiator. Target authentication provides security only for a specific iSCSI target. Multiple iSCSI initiators can use the same CHAP secret when connecting or each iSCSI initiator can have its own. The CHAP user name cannot include certain special characters.

Arrays that use the Fibre Channel protocol do not require CHAP user accounts.

chapuser --help

chapuser --list

chapuser --info *user_name*

chapuser --create *user_name*

--password *password*
[--description *text*]

chapuser --delete *user_name*
[--force]

chapuser --edit *user_name*
[--name *new_name*]
[--description *text*]
[--password *password*]

chapuser --add_initators *user_name*
[--initiator_name *iqn*]

chapuser --remove_initator *user_name*
[initiator_name *iqn*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all configured CHAP users for this array.	Operator
--info	*user_name*	Show information about a specific CHAP user.	Operator
--create	*user_name*	Create a new CHAP user and shared password for the CHAP user, who then gets assigned to an access control list (ACL)-associated volume. CHAP user authentication requires a CHAP secret that agrees as the challenge response.	Operator
	--password *password*	The CHAP secret. The CHAP secret must be must be between 12 and 16 alphanumeric characters, with no spaces or special characters, including ' " ` ~ ! @ # $ ^ & ( ) + [ ] {} * ; : ' " . , | < > ? / \ = % .	
	[--description *text*]	Plain-text description of the CHAP user. If there are spaces in the description, enclose the text in quotation marks.	
--delete	*user_name*	Delete the specified CHAP user. It deletes the CHAP user connection from the iSCSI initiator. If the CHAP user is assigned to any access control records, you must use the --force to delete.	Operator
	[--force]	Forcibly delete the specified CHAP user even if that user has active connections. It disregards any ACL records that have been applied.	

Option	Arguments and Suboptions	Description	Role
--edit	*user_name* [--name *new_name*]	Modify the name of the CHAP user. Changing the CHAP user name requires reconfiguring the iSCSI initiators to use the new CHAP user name.	Operator
	user_name [--description *text*]	Modify the description of the CHAP user. If there are spaces in the description, enclose the text in quotation marks.	
	user_name [--password *password*]	Modify the CHAP secret (shared password) used by the a CHAP user. This value must be between 12 and 16 alphanumeric characters, with no spaces or special characters, including ' " ` ~ ! @ # $ ^ & ( ) + [ ] {} * ; : ' " . , | < > ? / \ = %	
--add_initiators	*user_name*	Use --add_initiators to associate the specified CHAP user to specific iSCSI initiators. iSCSI Group Target uses this information to authenticate CHAP access to iSCSI initiators.	Operator
	[--initiator_name*iqn*]	Associate this CHAP user to the specified iSCSI initiators. If any specified initiator is already associated with another CHAP user, that user will be replaced by this CHAP user for future CHAP authentication.	
--remove_initiator	*user_name*	Use --remove_initiator to disassociate the CHAP user from a specific iSCSI initiator. iSCSI Group Target will not do CHAP authentication to this iSCSI initiator.	
	[--initiator_name*iqn*]	Disassociate the CHAP user from the specified iSCSI initiator. iSCSI Group Target will not do CHAP authentication to this iSCSI initiator.	

Examples

This example lists CHAP users in the array. In this example, two CHAP users exist.

```
Nimble OS $ chapuser --list
------------------------------+------------------------------------------
CHAP User Name                 Password
------------------------------+------------------------------------------
nimblechap                     nimblechapuser
guestchap                      minimalaccess
```

This example shows information about a CHAP user named "nimblechap".

```
Nimble OS $ chapuser --info nimblechap
Name: nimblechap
Description: generic chap user for the nimble array
Password: nimblechapuser
```

```
Created: Dec 17 2010 12:50:25
Last configuration change: Dec 17 2010 12:50:25
```

This example creates a CHAP user with the name "nimblestorage" and a valid password.

```
Nimble OS $ chapuser --create nimblestorage --description "nimble storage user"
 --password n1mblechapu5er
```

This example forcibly deletes a CHAP user named "storageuser".

```
Nimble OS $ chapuser --delete storageuser --force
```

This example modifies the name, description, and password for the CHAP user previously named "nimblechap".

Note: If you change the name or password of a CHAP user, it invalidates logins. Update the hosts with the change.

```
Nimble OS $ chapuser --edit nimblechap --name nimchap --description "nimble user"
 --password nimchapusr
```

ctrlr

The ctrlr command options display information about a controller on a Nimble array. Each array has an A and B controller (case-sensitive identifiers). One controller is always active and the other always takes the standby role. Controllers share IP addresses except for the support IP address. Each controller is also configured with the same set of NICs to provide redundancy in case of a failure. This configuration ensures that there is never a conflict during a failover. The IP addresses that are assigned to each NIC and interface remain the same. If a failover is triggered, then the new active controller has the same number of links to the switch as was the case before a failover occurred.

Synopsis

ctrlr --help

ctrlr --list
[--array *array_name*]

ctrlr --info *controller_name*
[--array *array_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List basic information about the controllers in the array where you are logged in.	Guest
	[--array *array_name*]	List basic information about all controllers on the specified array.	

Option	Arguments and Suboptions	Description	Role
--info	*controller_name*	List detailed information about the specified controller. Valid controller names are A and B, which are case sensitive.	Guest
	[--array *array_name*]	List detailed information about the specified controller on the specified array.	

Examples

This example lists basic information about the controllers on an array named *greyhound*.

```
Nimble OS $ ctrlr --list --array greyhound
----+----------+--------------------+----------------+--------+-------+--------
Name State      Hostname             Support IP       Power    Fans    Temper-

                                                      Supplies Status  -ature

                                                      Status           Status
----+----------+--------------------+----------------+--------+-------+--------
A    standby    greyhound-A          192.0.2.54       alerted  OK      OK
B    active     greyhound-B          192.0.2.55       alerted  OK      OK
```

This example shows information about controller B on an array named *greyhound*.

```
Nimble OS $ ctrlr --info B --array greyhound
Name: B
Serial number: AC-102724-C2
State: active
Hostname: greyhound-B
Support IP address: 192.0.2.55
Support IP netmask: 255.255.255.0
Support IP nic: eth1
Hostname:
Power supply: alerted
        power-supply1 at left rear: ok
        power-supply2 at right rear: missing
Cooling fans: OK
        fan1 at lower front of controller B: ok, speed: 10608rpm
        fan2 at lower left rear of controller B: ok, speed: 10608rpm
        fan3 at lower right rear of controller B: ok, speed: 10608rpm
        fan4 at upper right front of controller B: ok, speed: 8295rpm
        fan5 at upper left front of controller B: ok, speed: 8910rpm
        fan6 at upper left rear of controller B: ok, speed: 8520rpm
Temperature sensors: OK
        motherboard at motherboard: ok, temperature: 22C
        bp-temp2 at right-side backplane: ok, temperature: 34C
System partition status: OK
Last AutoSupport contact: N/A
```

date

The date command options display or modify the date and time of the array. If the array is connected to a valid NTP server IP address, then the time settings on the array synchronize with the time settings on the NTP server. In this case, do not modify the date locally. However, if the array is connected to an invalid NTP server IP address, then the time settings on the array do not synchronize with the time settings on the NTP server.

Synopsis

date --help

date
[--utc]
[--edit *{hh:mm[:ss]|'YYYY-MM-DD hh:mm[:ss]'*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--utc	N/A	List the Coordinated Universal Time on an array.	Guest
--edit	*{hh:mm[:ss]|'YYYY-MM-DD hh:mm[:ss]"*	Use --edit with care. Modify the date and time on an array and disable the connection between the array and a previously specified NTP server. If specified with --utc, modify Coordinated Universal Time to the supplied argument.	Power User

Examples

This example lists the date and time on an array as Pacific Daylight Time.

```
Nimble OS $ date
Wed Jul  9 15:29:59 PDT 2014
```

This example lists the date and time on an array as Coordinated Universal Time.

```
Nimble OS $ date --utc
Wed Jul  9 22:29:37 UTC 2014
```

This example modifies the date and time on an array.

```
Nimble OS $ date --edit '2013-07-08 15:33:27'
WARNING: Setting date manually disables the NTP server setting
```

disk

The disk command options manage disks on an array. An array supports hard disk drives (HDDs) and solid state drives (SSDs). The size of the disk depends on the array model. Some HDDs are used to build the RAID and others can be spares that are used to replace failing drives. SSDs can only be active, and they are not part of the RAID configuration.

When you remove a hard disk drive (HDD) from an array, it is either removed from the existing RAID or replaced as a spare, depending on the state of the RAID. If you remove a solid state disk (SSD), it is removed from the existing cache capacity.

Synopsis

disk --help

disk --list
[--array *array_name*]

[--shelf_location *shelf_location*]

disk --info *slot*
[--array *array_name*]
[--shelf_location *shelf_location*]

disk --add *slot*
[--array *array_name*]
--shelf_location *shelf_location*
[--force]

disk --remove *slot*
[--array *name*]
--shelf_location*shelf_location*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List basic information about the disks on an array, showing the type, slot, serial number, size, disk state, RAID status, shelf serial number, and shelf location for each.	Guest
	[--array *array_name*]	List basic information about the disks on the specified array. Specify this option in a multi-array configuration. It is optional in a single-array configuration.	
	[--shelf_location *shelf_location*]	List basic information about the disks on the specified shelf. If this option is not specified, then all disks on all shelves are listed.	
--info	*slot*	List detailed information about the disk at a specified slot number in a single-array configuration.	Guest
	[--array *array_name*]	List detailed information about all disks for a specified array in a multi- array configuration. Include this option in a multi-array configuration. It is optional in a single-array configuration.	
	[--shelf_location *shelf_location*]	List detailed information about the disk at the specified shelf location. If this option is not specified, then only disks for the head shelf A (A.0) or B (B.0) are listed.	

Option	Arguments and Suboptions	Description	Role
--add	*slot*	Add a disk at the specified slot on a single-array configuration.	Power User
	[--array *array_name*]	Add a disk at the specified slot on a multi-array configuration. Include this option in a multi-array configuration. It is optional in a single-array configuration.	
	--shelf_location *shelf_location*	Add a disk at the specified shelf location on a specified array. This option is mandatory.	
	[--force]	Forcibly add the specified disk. This option is mandatory if the disk to add is in a foreign state.	
--remove	*slot*	Remove a disk from the specified slot in a single-array configuration.	Power User
	[--array *array_name*]	Remove a disk from the specified slot on a specified array in a multi-array configuration. Include this option in a multi-array configuration. It is optional in a single-array configuration.	
	--shelf_location*shelf_location*	Remove a disk from the specified slot on the specified shelf location. This option is mandatory.	

Examples

This example lists all disks on a standalone (single) array.

```
Nimble OS $ disk --list
------+--------------------+----+---------+-------+--------+---------+-----
Slot # Serial #             Type Disk Size Disk    RAID     Shelf     Shelf
                                 (GB)      State   Status   Serial    Loca-
                                                                      -tion
------+--------------------+----+---------+-------+--------+---------+-----
     1 WD-WCAW32373319      HDD    1000.20 in use  spare    AA-100257 A.0
     2 WD-WCAW30555873      HDD    1000.20 in use  okay     AA-100257 A.0
     3 WD-WMAW30010098      HDD    1000.20 in use  okay     AA-100257 A.0
     4 WD-WCAW30549295      HDD    1000.20 in use  okay     AA-100257 A.0
     5 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     6 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     7 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     8 CVPO040402ZP080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
     9 CVPO04030150080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
    10 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    11 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    12 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    13 WD-WCAW30597086      HDD    1000.20 in use  okay     AA-100257 A.0
    14 WD-WCAW30606284      HDD    1000.20 in use  okay     AA-100257 A.0
    15 WD-WCAW30606703      HDD    1000.20 in use  okay     AA-100257 A.0
    16 WD-WCAW30527176      HDD    1000.20 in use  okay     AA-100257 A.0
     1 CVPR208103WZ600FGN   SSD     600.13 in use  N/A      AC-100182 A.1
     2 Z29036W600009136H45Q HDD    3000.59 in use  okay     AC-100182 A.1
     3 Z290LRCK00009136H440 HDD    3000.59 in use  spare    AC-100182 A.1
     4 9XK0H5N600009135E6GN HDD    3000.59 in use  okay     AC-100182 A.1
```

```
     5 9XK06BFJ0000S111K944 HDD    3000.59 in use  okay     AC-100182 A.1
     6 Z290LAW80000S110DX86 HDD    3000.59 in use  okay     AC-100182 A.1
     7 Z290372V0000914024J2 HDD    3000.59 in use  okay     AC-100182 A.1
     8 9XK0KBZ400009131N1NV HDD    3000.59 in use  okay     AC-100182 A.1
     9 Z2909KED00009133BHZ5 HDD    3000.59 in use  okay     AC-100182 A.1
    10 Z290KYG600009138MTYB HDD    3000.59 in use  okay     AC-100182 A.1
    11 Z290LNP9000091380XUM HDD    3000.59 in use  okay     AC-100182 A.1
    12 9XK06NF00000S111K9X6 HDD    3000.59 in use  okay     AC-100182 A.1
    13 9XK07TQN0000S112MQAQ HDD    3000.59 in use  okay     AC-100182 A.1
    14 9XK06B6G0000S111K9MH HDD    3000.59 in use  okay     AC-100182 A.1
    15 9XK06AYR0000S110UGRY HDD    3000.59 in use  okay     AC-100182 A.1
    16 9XK06NEP0000S110UGCD HDD    3000.59 in use  okay     AC-100182 A.1
```

This example lists only the disks on the specified shelf.

```
Nimble OS $ disk --list --shelf_location A.0
------+--------------------+----+---------+-------+--------+---------+-----
Slot # Serial #             Type Disk Size Disk    RAID     Shelf     Shelf
                                 (GB)      State   Status   Serial    Loca-
                                                                      -tion
------+--------------------+----+---------+-------+--------+---------+-----
     1 WD-WCAW32373319      HDD    1000.20 in use  spare    AA-100257 A.0
     2 WD-WCAW30555873      HDD    1000.20 in use  okay     AA-100257 A.0
     3 WD-WMAW30010098      HDD    1000.20 in use  okay     AA-100257 A.0
     4 WD-WCAW30549295      HDD    1000.20 in use  okay     AA-100257 A.0
     5 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     6 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     7 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     8 CVPO040402ZP080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
     9 CVPO04030150080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
    10 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    11 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    12 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    13 WD-WCAW30597086      HDD    1000.20 in use  okay     AA-100257 A.0
    14 WD-WCAW30606284      HDD    1000.20 in use  okay     AA-100257 A.0
    15 WD-WCAW30606703      HDD    1000.20 in use  okay     AA-100257 A.0
    16 WD-WCAW30527176      HDD    1000.20 in use  okay     AA-100257 A.0
```

This example shows information about the hard disk drive in slot 1 in a single-array configuration. The output is truncated to avoid line wrapping.

```
Nimble OS $ disk --info 1
Shelf location #: A.0
Slot #: 1
HBA: 0
Port: 15
Serial: WD-WCAW32373319
Type: HDD
Disk size: 1000.2 GB
Disk state: in use
Raid state: spare
Percent resynchronized: N/A
Current resync speed: N/A
Average resync speed: N/A
Model: WDC WD1003FBYX-0
Vendor: Nimble
Firmware version: 1V01
Smart Data Attributes:
Raw read error rate(value/trough/threshold/raw/updated): 200/200/51/5/Ju...
Spin up time(value/trough/threshold/raw/updated): 174/172/21/4291/Jul 15...
Start stop count(value/trough/threshold/raw/updated): 100/100/0/393/Jul ...
```

```
Reallocated sector count(value/trough/threshold/raw/updated): 200/200/1 ...
Seek error rate(value/trough/threshold/raw/updated): 200/200/0/0/Jul 15 ...
Power on hours(value/trough/threshold/raw/updated): 91/91/0/6890/Jul 15 ...
Power cycle count(value/trough/threshold/raw/updated): 100/100/0/392/Jul...
Power off retract count(value/trough/threshold/raw/updated): 200/200/0/3...
Temperature C(value/trough/threshold/raw/updated): 116/99/0/31/Jul 15 20...
Current pending sector(value/trough/threshold/raw/updated): 200/200/0/0/...
Offline uncorrectable(value/trough/threshold/raw/updated): 200/200/0/0/ ...
UDMA CRC error count(value/trough/threshold/raw/updated): 200/200/0/0/Ju...
```

This example shows information about the disk in slot 1. Because the --shelf_location option is specified in this example, only the disks that are located in slot 1 at the specified shelf location are listed. The output is truncated to avoid line wrapping.

```
Nimble OS $ disk --info 1 --shelf_location A.1
Shelf location #: A.1
Slot #: 1
HBA: 0
Port: 15
Serial: CVPR208103WZ600FGN
Type: SSD
Disk size: 600.127 GB
Disk state: in use
Raid state: N/A
Percent resynchronized: N/A
Current resync speed: N/A
Average resync speed: N/A
Model: INTEL SSDSA2CW60
Vendor: Nimble
Firmware version: 0362
Smart Data Attributes:
Spin up time(value/trough/threshold/raw/updated): 100/100/0/0/Jul 15 201...
Start stop count(value/trough/threshold/raw/updated): 100/100/0/0/Jul 15...
Reallocated sector count(value/trough/threshold/raw/updated): 100/100/0/...
Power on hours(value/trough/threshold/raw/updated): 100/100/0/10495/Jul ...
Power cycle count(value/trough/threshold/raw/updated): 100/100/0/27/Jul ...
Power off retract count(value/trough/threshold/raw/updated): 100/100/0/2...
Host write count(value/trough/threshold/raw/updated): 100/100/0/344912/J...
Available reserved space(value/trough/threshold/raw/updated): 100/100/10...
Media wearout indicator(value/trough/threshold/raw/updated): 99/99/0/0/J...
```

This example adds a disk to slot 13 at the specified shelf location A.0 and verifies the change in a single-array configuration. The disk resynchronizes with the rest of the array. The --shelf_location option is mandatory in this example.

```
Nimble OS $ disk --add 13 --shelf_location A.0
Nimble OS $ disk --list
------+--------------------+----+---------+-------+--------+---------+-----
Slot # Serial #             Type Disk Size Disk    RAID     Shelf     Shelf
                                 (GB)      State   Status   Serial    Loca-
                                                                      -tion
------+--------------------+----+---------+-------+--------+---------+-----
     1 WD-WCAW32373319      HDD    1000.20 in use  okay     AA-100257 A.0
     2 WD-WCAW30555873      HDD    1000.20 in use  okay     AA-100257 A.0
     3 WD-WMAW30010098      HDD    1000.20 in use  okay     AA-100257 A.0
     4 WD-WCAW30549295      HDD    1000.20 in use  okay     AA-100257 A.0
     5 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     6 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     7 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     8 CVPO040402ZP080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
     9 CVPO04030150080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
    10 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    11 N/A                  N/A        N/A void    N/A      AA-100257 A.0
```

```
    12 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    13 WD-WCAW30597086      HDD    1000.20 in use  spare    AA-100257 A.0
    14 WD-WCAW30606284      HDD    1000.20 in use  okay     AA-100257 A.0
    15 WD-WCAW30606703      HDD    1000.20 in use  okay     AA-100257 A.0
    16 WD-WCAW30527176      HDD    1000.20 in use  okay     AA-100257 A.0
```

This example removes a disk from slot 13 at the specified shelf location A.0 and verifies the change in a single-array configuration. The --shelf_location option is mandatory in this example.

```
Nimble OS $ disk --remove 13 --shelf_location A.0
Nimble OS $ disk --list
------+--------------------+----+---------+-------+------------------+-----
Slot # Serial #             Type Disk Size Disk    RAID     Shelf     Shelf
                                 (GB)      State   Status   Serial    Loca-
                                                                      -tion
------+--------------------+----+---------+-------+--------+---------+-----
     1 WD-WCAW32373319      HDD    1000.20 in use  resynch  AA-100257 A.0
     2 WD-WCAW30555873      HDD    1000.20 in use  okay     AA-100257 A.0
     3 WD-WMAW30010098      HDD    1000.20 in use  okay     AA-100257 A.0
     4 WD-WCAW30549295      HDD    1000.20 in use  okay     AA-100257 A.0
     5 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     6 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     7 N/A                  N/A        N/A void    N/A      AA-100257 A.0
     8 CVPO040402ZP080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
     9 CVPO04030150080JGN   SSD      80.03 in use  N/A      AA-100257 A.0
    10 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    11 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    12 N/A                  N/A        N/A void    N/A      AA-100257 A.0
    13 WD-WCAW30597086      HDD    1000.20 removed N/A      AA-100257 A.0
    14 WD-WCAW30606284      HDD    1000.20 in use  okay     AA-100257 A.0
    15 WD-WCAW30606703      HDD    1000.20 in use  okay     AA-100257 A.0
    16 WD-WCAW30527176      HDD    1000.20 in use  okay     AA-100257 A.0
     1 CVPR208103WZ600FGN   SSD     600.13 in use  N/A      AC-100182 A.1
     2 Z29036W600009136H45Q HDD    3000.59 in use  okay     AC-100182 A.1
     3 Z290LRCK00009136H440 HDD    3000.59 in use  spare    AC-100182 A.1
     4 9XK0H5N600009135E6GN HDD    3000.59 in use  okay     AC-100182 A.1
     5 9XK06BFJ0000S111K944 HDD    3000.59 in use  okay     AC-100182 A.1
     6 Z290LAW80000S110DX86 HDD    3000.59 in use  okay     AC-100182 A.1
     7 Z290372V0000914024J2 HDD    3000.59 in use  okay     AC-100182 A.1
     8 9XK0KBZ400009131N1NV HDD    3000.59 in use  okay     AC-100182 A.1
     9 Z2909KED00009133BHZ5 HDD    3000.59 in use  okay     AC-100182 A.1
    10 Z290KYG600009138MTYB HDD    3000.59 in use  okay     AC-100182 A.1
    11 Z290LNP9000091380XUM HDD    3000.59 in use  okay     AC-100182 A.1
    12 9XK06NF00000S111K9X6 HDD    3000.59 in use  okay     AC-100182 A.1
    13 9XK07TQN0000S112MQAQ HDD    3000.59 in use  okay     AC-100182 A.1
    14 9XK06B6G0000S111K9MH HDD    3000.59 in use  okay     AC-100182 A.1
    15 9XK06AYR0000S110UGRY HDD    3000.59 in use  okay     AC-100182 A.1
    16 9XK06NEP0000S110UGCD HDD    3000.59 in use  okay     AC-100182 A.1
```

encryptkey

The encryptkey command options manage the encryption master key and its associated passphrase.

The master key is protected by a passphrase and is used to encrypt the keys that are used to encrypt volume data. This command provides options to create, enable, disable, delete, and purge the master key, and to change the passphrase that is used to protect the master key.

Data encryption keys for volumes are encrypted by using a master key that must be initialized before encrypted volumes can be created. The master key in turn is protected by a passphrase that is set when the master key is created. The passphrase may have to be entered to enable the master key when it is not available, for example, after an array reboot.

After creating a master key, it can be in two possible states: active and inactive. In the active state, the master key is available to unlock the keys that are used to encrypt volume data. When inactive, the master key is not available, and access to encrypted data is not possible. To change the state of the master key from inactive to active, use --enable_master and enter the passphrase.

> **Note:** No characters appear on the screen as you enter the passphrase for any of the encryptkey command options, not even asterisks to indicate an obfuscated value.

Synopsis

encryptkey --help

encryptkey --info

encryptkey --create_master

encryptkey --enable_master

encryptkey --disable_master

encryptkey --delete_master

encryptkey --change_passphrase

encryptkey --purge_inactive

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--info	N/A	Display the status of the master key. Initialized is true only if the master key has been created. Active is true only if the master key has been unlocked using the passphrase and is available for use.	Operator
--create_master	N/A	Create the passphrase-encrypted master key. You must interactively specify and confirm the passphrase. The passphrase value can be any printable characters, and its length must be between 8 and 64 characters, inclusive. The group encryption_scope is set to group and the group encryption_cipher is set to aes-256-xts.	Administrator
--enable_master	N/A	Activate the master key. You must interactively specify the correct passphrase.	Operator

Option	Arguments and Suboptions	Description	Role
--disable_master	N/A	Deactivate the master key and disable encryption. You mus interactively specify the correct passphrase. **Note:** You can still take snapshots of encrypted volumes after disabling the master key.	Administrator
--delete_master	N/A	Delete the master key. All encrypted volumes must first be deleted. This allows the recovery of a group when the passphrase has been lost. The encryption scope and cipher are both reset to none.	Administrator
--change_passphrase	N/A	Decrypt the master key and verify using the old passphrase, and re-encrypt using a new passphrase.	Administrator
--purge_inactive	--age *age_in_hours*	Purge encryption keys that have been inactive for the specified number of hours or longer. If you do not specify an age in hours, inactive keys will be purged immediately. The *age_in_hours* is the minimum number of hours keys may be inactive before they are purged. To purge inactive keys immediately, enter 0.	Administrator

Examples

This example identifies that encryption is uninitialized and inactive on the Nimble Storage group.

```
Nimble OS $ encryptkey --info
Initialized: No
Active: No
```

This example shows detailed information about a Nimble Storage group where encryption is enabled (initialized and active).

```
Nimble OS $ encryptkey --info
Initialized: Yes
Active: Yes
```

Note: If you disable the master key by using the --disable_master option, then encryption is inactive in the --info output.

This example creates the master key and specifies its passphrase interactively. This command also enables the master key.

```
Nimble OS $ encryptkey --create_master
Enter new passphrase:
Retype new passphrase:
```

This example re-enables the master key after the --disable_master option is used.

```
Nimble OS $ encryptkey --enable_master
Enter passphrase:
```

This example disables the master key after creating it. Encrypted volumes go offline and are inoperative until you enable the master key. If you have volumes that are not encrypted, then those volumes are still accessible. The --disable_master capability is provided only through the CLI, not the GUI.

```
Nimble OS $ encryptkey --disable_master
Enter passphrase:
```

This example deletes the master key. This allows recovering a system where the passphrase has been lost. If this happens, any encrypted volumes are permanently inaccessible, and they must be deleted before this operation will succeed. The --delete_master capability is provided only through the CLI, not the GUI.

```
Nimble OS $ encryptkey --delete_master
```

This example changes the passphrase for the master key interactively. There is no indication of successful completion. If you type an incorrect value for the current passphrase, or if you type and confirm values that do not match for the new passphrase, then an error identifies the issue.

```
Nimble OS $ encryptkey --change_passphrase
Enter passphrase:
Enter new passphrase:
Retype new passphrase:
```

This example configures the number of hours after which to purge inactive keys to two hours.

```
Nimble OS $ encryptkey --purge_inactive --age 2
```

failover

The failover command options perform a controller failover on an array. The active controller restarts and assigns the active role to the standby controller.

Synopsis

failover --help

failover
[--array *array_name*]
[--force]
[--non_interactive]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--array	*array_name*	Perform a controller failover on the specified array, which must be running in redundant controller mode (active/standby). You must specify the --array *array_name* option in a multi-array configuration. Specifying this option is not required in a single-array configuration.	Power User
[--force]	N/A	Forcibly fail over the controllers on an array even if the validation checks that the system runs before initiating the failover fail.	
[--non_interactive]	N/A	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly taking the array offline.	

Examples

This example forces a failover of a controller in a single-array configuration.

```
Nimble OS $ failover
WARNING: This operation will fail over the array.
Type yes to continue, no to cancel: yes
INFO: Initiated controller failover.
```

This example forces a failover of a controller named greyhound in a multi-array configuration.

```
Nimble OS $ failover --array greyhound
WARNING: This operation will fail over the array.
Type yes to continue, no to cancel: yes
INFO: Initiated controller failover.
```

fc

The fc command is used to manage Fibre Channel interfaces.

Synopsis

fc --help

fc --list
[--array {*name|serial*}]

fc --info *interface_name*
[--array {*name|serial*}]
[--ctrlr {*A|B*}]

fc --ping *interface_name*
[--wwpn *wwpn*]
[--fcid *fcid*]
[--count *count*]

fc --edit *interface_name*
[--array {*name|serial*}]
[--ctrlr {*A|B*}]
[--admin_state {*online|offline*}]
[--force]

fc --update_config

fc --regenerate_wwn *wwnn_base*
[--non_interactive]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List all discovered Fibre Channel interfaces for arrays in the group.	Guest
	[--array {*name|serial*}]	List all discovered Fibre Channel interfaces for the specified array.	
--info	*interface_name*	Show detailed information about the specified Fibre Channel interface.	Guest
	[--array {*name|serial*}]	Show detailed information about a Fibre Channel interface on a specified array. Provide the array name or serial number.	
	[--ctrlr {*A|B*}]	Show detailed information about a Fibre Channel interface on a specified controller. Valid controller names are A and B, which are case sensitive.	
--ping	*interface_name*	Send ECHO ELS frame to specified Worldwide Port Number (WWPN) or port FC-ID from Fibre Channel interface on local controller.	
	[--wwpn *wwpn*]	Specify the remote WWPN.	Power User
	[--fcid *fcid*]	Specify the remote port FC-ID in six hexadecimal digits format.	
	[--count *count*]	Repeat ping for the specified number of times.	

Option	Arguments and Suboptions	Description	Role
--edit	*interface_name*	Modify attributes of the specified Fibre Channel interface.	Power User
	[--array {*name*|*serial*}]	Modify a Fibre Channel interface on the specified array.	
	[--ctrlr {*A*|*B*}]	Modify a Fibre Channel interface on the specified controller. Valid controller names are A and B, which are case sensitive.	
	[--admin_state {*online*|*offline*}]	Set the desired administrative state for Fibre Channel interface. Valid states are *online* and *offline*. Setting a Fibre Channel interface offline with connected initiators will fail, unless the command --force option is used.	
	[--force]	Forcibly modify the administrative state of the specified Fiber Channel interface. This option is required when setting a Fibre Channel interface to the *offline* state with connected initiators.	
--update_config	N/A	Use the fc --update_config command only if instructed to do so by Nimble Storage Support. Update the Fibre Channel configuration after a hardware upgrade or downgrade.	Power User
--regenerate_wwn	*wwnn_base*	Regenerate the World Wide Node Name (WWNN) and all World Wide Port Names (WWPNs) of Fibre Channel interfaces in the group to set the non-fixed bits of the WWNN to the specified *wwnn_base*. The value is six hexadecimal digits between 00:00:01 and FF:FF:FF, inclusive. For example, if 12:34:56 is specified, the WWNN will be set to 56:c9:ce:90:12:34:56:00 and the first WWPN will be 56:c9:ce:90:12:34:56:01. The WWPNs will be numbered in order from lowest interface to highest interface (e.g., fc1.1, fc2.1, fc5.1, etc.), starting first with controller A then controller B. All Fibre Channel interfaces must be both administratively and operationally offline before proceeding. After the regenerate completes, all interfaces will be brought back online.	Power User
	[--non_interactive]	Do not prompt for confirmation.	

Examples

This example lists the Fibre Channel Interfaces on a CS300 array named *greyhound*. The output is truncated to avoid line wraps.

```
Nimble OS $ fc --list --array greyhound
Array: greyhound
------+-----+-----+------+----+------------------------+----------------------
Name  Ctrlr Admin Fabric Link WWNN                       WWPN
           Status
------+-----+-----+------+----+------------------------+----------------------
fc5.1  A    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:01
fc6.1  A    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:02
fc5.1  B    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:03
fc6.1  B    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:04
```

This example shows information about the fc5.1 Fibre Channel interface on controller A of an array named *greyhound*.

```
Nimble OS $ fc --info fc5.1 --ctrlr A
Name: fc5.1
Array: rtp-array33
Controller: A
Admin Status: online
Link Status: up
Operational Status: operational
Fabric Logged In: yes
Link Speed: 8G
Max Link Speed: 16G
WWNN: 56:c9:ce:90:ac:d1:04:00
WWPN: 56:c9:ce:90:ac:d1:04:01
FC-ID: 0xe007d0
PCI bus ID: 0000:84:00.0
HBA port: 0
HBA slot: 2
Firmware version: 11.2.156.27
Fabric switch name: rtp-b10-nex-1
Fabric switch port: fc2/9
Fabric switch WWNN: 20:64:8c:60:4f:bb:cb:01
Fabric switch WWPN: 20:49:8c:60:4f:bb:cb:00
Fabric WWN: 20:64:00:de:fb:2c:50:c1
SFP Rx Power: 577 uW
SFP Tx Power: 606 uW

Connected Initiators:
-----------------+-----------------------+-----------------------
Alias              WWNN                    WWPN
-----------------+-----------------------+-----------------------
c22_esx1_hba1      20:00:00:90:fa:53:a0:a0 10:00:00:90:fa:53:a0:a0
c22_esx1_hba2      20:00:00:90:fa:53:a0:a1 10:00:00:90:fa:53:a0:a1
c22_win7_port0     20:00:00:90:fa:53:9f:08 10:00:00:90:fa:53:9f:08
c22_win7_port1     20:00:00:90:fa:53:9f:09 10:00:00:90:fa:53:9f:09
c22win8hba1        20:00:00:90:fa:53:43:8c 10:00:00:90:fa:53:43:8c
c22win8hba2        20:00:00:90:fa:53:43:8d 10:00:00:90:fa:53:43:8d
c23_esx2_hba1      20:00:00:90:fa:53:a2:80 10:00:00:90:fa:53:a2:80
c23_esx2_hba2      20:00:00:90:fa:53:a2:81 10:00:00:90:fa:53:a2:81
c23_esx3_hba2      20:00:00:0e:1e:19:5e:a0 21:00:00:0e:1e:19:5e:a0
c23_esx3_hba3      20:00:00:0e:1e:19:5e:a1 21:00:00:0e:1e:19:5e:a1
```

This example sets the fc5.1 Fibre Channel interface on controller A to be offline.

```
Nimble OS $ fc --edit fc5.1
--controller A --admin_state
offline --force
```

This example updates the Fibre Channel configuration upon the advice of HPE Nimble Storage Support after performing a hardware upgrade or downgrade (adding or removing Fibre Channel HBAs) on both controllers of an array.

```
Nimble OS $ fc --update_config
```

This example allows Fibre Channel zoning to be configured before the HPE Nimble Storage array you purchased arrives on site. Follow these steps:

1. Read this fc command section to review the process to follow.
2. Plan zoning for the Nimble array you purchased, with the understanding that you can change the *wwnn_base* value after you set up the array.
3. In a standard maintenance window, configure zoning for the HPE Nimble Storage array on the Fibre Channel switches.
4. When the HPE Nimble Storage array arrives on site, install and set up the array by following the instructions in the *Installation Guide* that applies to your array model. This step applies unique system-assigned WWNN and WWPN values, which you will modify later in this procedure.
5. Complete the commands in the following example, using values that are appropriate for your environment.

The example shows the remaining steps in the process on an HPE Nimble Storage array named *greyhound*, where the interfaces are named fc5.1 and fc6.1.

```
Nimble OS $ fc --list --array greyhound
Array: greyhound
------+-----+-----+------+----+-----------------------+----------------------
Name  Ctrlr Admin Fabric Link WWNN                     WWPN
            Status
------+-----+-----+------+----+-----------------------+----------------------
fc5.1  A    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:01
fc6.1  A    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:02
fc5.1  B    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:03
fc6.1  B    online  yes  16G  56:c9:ce:90:ad:d6:3a:00  56:c9:ce:90:ad:d6:3a:04

Nimble OS $ fc --edit fc5.1 --ctrlr A --admin_state offline

Nimble OS $ fc --edit fc6.1 --ctrlr A --admin_state offline

Nimble OS $ fc --edit fc5.1 --ctrlr B --admin_state offline

Nimble OS $ fc --edit fc6.1 --ctrlr B --admin_state offline

Nimble OS $ fc --list --array greyhound
...
...
Nimble OS $ fc --regenerate_wwn 12:34:56 --non_interactive

Nimble OS $ fc --list --array greyhound
Array: greyhound
------+-----+-----+------+----+-----------------------+----------------------
Name  Ctrlr Admin Fabric Link WWNN                     WWPN
            Status
------+-----+-----+------+----+-----------------------+----------------------
fc5.1  A    online  yes  16G  56:c9:ce:90:12:34:56:00  56:c9:ce:90:12:34:56:01
fc6.1  A    online  yes  16G  56:c9:ce:90:12:34:56:00  56:c9:ce:90:12:34:56:02
fc5.1  B    online  yes  16G  56:c9:ce:90:12:34:56:00  56:c9:ce:90:12:34:56:03
```

```
fc6.1  B    online  yes  16G  56:c9:ce:90:12:34:56:00  56:c9:ce:90:12:34:56:04
```

In this example, because 12:34:56 is specified as the *wwnn_base* value, the WWNN is changed from 56:c9:ce:90:ad:d6:3a:00 to 56:c9:ce:90:12:34:56:00 and the first WWPN is set to 56:c9:ce:90:12:34:56:01. The WWPNs are numbered in order from the lowest interface to the highest interface, starting first with controller A and then controller B, as shown in the final fc --list output in the previous command examples. The command with the --regenerate_wwn option brings the Fibre Channel interfaces back online automatically.

folder

The folder command options allow you to manage folders. Folders are a way to group volumes, and can be used to provision storage for third party plugins such as VVOLs and SMI-S. In the case of VVOLs, folders appear as Storage Containers. In the case of SMI-S, folders appear as Storage Pools.

All folder commands that accept a volume name as an argument will accept an optional --pool argument to differentiate those folders that have the same name but reside in different pools.

Synopsis

folder --help

folder --list
[--pool *pool_name*]
[--access_protocol {*iscsi*|*fc*}]

folder --create *folder_name*
[--pool *pool_name*]
[--description *text*]
[--usage_limit *mebibytes*]
[--provisioned_limit *mebibytes*]
[--overdraft_limit_pct *percent*]
[--iops_limit *iops*]
[--mbps_limit *mbps*]
[--agent_type {*none*|*smis*|*vvol*|*openstack*}]
[--appserver *vcenter_name*]
[--perfpolicy *perfpolicy*]
[--access_protocol {*iscsi*|*fc*}]

folder --edit *folder_name*
[--pool *pool_name*]
[--name *new_name*]
[--pool *pool_name*]
[--description *text*]
[--usage_limit *mebibytes*]
[--provisioned_limit *mebibytes*]
[--overdraft_limit_pct *percent*]
[--iops_limit *iops*]
[--mbps_limit *mbps*]
[--agent_type {*none*|*smis*|*vvol*|*openstack*}]
[--appserver *vcenter_name*]
[--force]
[--perfpolicy *perfpolicy*]

folder --delete *folder_name*

[--pool *pool_name*]
[--force]

folder --info *folder_name*
[--pool *pool_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List folders in a group.	Operator
	[--pool *pool_name*]	Limit the list of folders to the specified pool.	
	[--access_protocol {*iscsi*|*fc*]	Limit the list of folders to the specified access protocol. This option only applies when folders have an agent type of "vvol".	

Option	Arguments and Suboptions	Description	Role
--create	*folder_name*	Create a folder.	Operator
	[--pool *pool_name*]	Name of the pool to contain the folder. If the pool option is not specified, the folder will be created in the default pool.	
	[--description *text*]	Description of folder.	
	[--usage_limit *mebibytes*]	Limit for the folder in mebibytes. To set the folder size to unlimited (size limited by parent pool), use "unlimited". Default: unlimited folder size is equal to the pool capacity.	
	[--provisioned_limit *mebibytes*]	Provisioned limit for the folder in mebibytes. To set the folder provisioned limit to unlimited, use "unlimited".	
	[--overdraft_limit_pct *percent*]	Overdraft percentage by which the folder can exceed the usage limit. The default is 0 percent. Overdraft percent limit requires that the --usage_limit be specified.	
	[agent_type {*none*|*smis*|*vvol*|*openstack*}]	External management agent type. Choices are "none", "smis", "vvol", or "openstack". Default: "none".	
	[--appserver *vcenter_name*]	Name of the vCenter associated with this folder. This option only applies to folders with an agent_type of "vvol".	
	[--perfpolicy *perfpolicy*]	Default performance policy for volumes created in this folder. The default performance policy is "default". This option only applies to folders with agent_type "smis".	
	[--access_protocol {*iscsi*|*fc*}]	Access protocol for the folder. If this option is not specified, the access protocol supported by the group is used. If the group supports multiple protocols, the default will be Fibre Channel.	
	[--iops_limit *iops*]	The IOPS limit for the folder. To set the IOPS limit to unlimited, use "unlimited".	
	[--mbps_limit *mbps*]	Throughput limit for the folder in MiBs. To set the throughput limit to unlimited, use "unlimited".	

Option	Arguments and Suboptions	Description	Role
--edit	*folder_name*	Modify attributes of the specified folder.	Operator
	[--pool *pool_nam*]	Name of the pool that contains the folder. If the pool option is not specified, the default pool will be used.	
	[--name *new_name*]	Modify the name of the specified folder. This will correspondingly change the name of the Storage Container or the Storage Pool, if applicable.	
	[--description *text*]	Modify the description of the specified folder.	
	[--usage_limit *mebibytes*]	Modify the usage limit of the specified folder. To set the folder size to unlimited (size limited by parent pool), use "unlimited".	
	[--appserver *vcenter_name*]	Change the vCenter associated with this folder. This option only applies to folders with an agent_type "vvol" with no vCenter association.	
	[--usage_limit *mebibytes*]	Modify the usage limit of the specified folder. To set the folder usage limit to unlimited (size limited by parent pool), use "unlimited".	
	[--provisioned_limit *mebibytes*]	Modify the provisioned limit of the specified folder. To set the provisioned limit to unlimited (size limited by parent pool), use "unlimited".	
	[--overdraft_limit_pct *percent*]	Modify the overdraft limit percentage for the specified folder. The overdraft limit is a percentage of the folder usage limit.	
	[--force]	Forcibly associate the folder with the specified application server. This option is required when the folder already associates with an application server.	
	[--perfpolicy *perfpolicy*]	Modify the performance policy assigned by default to volumes created in this folder by the "create volume" API.	
	[--iops_limit *iops*]	Modify the IOPS limit for the folder. To set the IOPS limit to unlimited, use "unlimited".	
	[--mbps_limit *mbps*]	Modify the throughput limit for the specified folder. To set the throughput limit to unlimited, use "unlimited".	

Option	Arguments and Suboptions	Description	Role
--delete	*folder_name*	Delete the specified folder. No volumes should be present under this folder.	Operator
	[--pool *pool_name*]	Name of the pool that contains the folder. If the pool option is not specified, the default pool is used.	
	[--force]	Forcibly delete a folder managed by vCenter.	
--info	*folder_name*	List detailed information about the specified folder.	Operator
	[--pool *pool_name*]	Name of the pool that contains the folder. If the pool option is not specified, the default pool is used.	

Examples

This example lists the folders.

```
Nimble OS $ folder --list
----------------------+-----------------------+---------+----------------------
Path                        Description         Agent Type     Usage Limit(MiB)
----------------------+-----------------------+---------+----------------------
default:/Folder-1         Test                    none      Unlimited (3530971)
default:/Folder-Mal                                 smis                      12
```

This example lists detailed information about the folder named *Folder-Mal*

```
Nimble OS $ folder --info Folder-Mal
Path/URI: default:/Folder-Mal
Description:
Agent Type: smis
Application UUID: 00000002000040028ecb759c52142abb
Volume mapped usage (MiB): 0
Uncompressed snapshot usage including pending deletes (MiB): 0
Total volume and snapshot usage (MiB): 0
Overdraft Limit (%): 10
Usage Limit (MiB): 12
Provisioned Size (MiB): 2048
Provisioned Limit (MiB): Unlimited
Volume count: 1
Volumes in the folder: Vol-mal3
Snapshot count: deprecated
Snapshot collection count: deprecated
Volume compression: 1.00X
Unused reserve (MiB): 0
Snapshot compression: 1.00X
Performance policy: default
Application server: N/A
IOPS Limit: 456
Throughput Limit (MiB/s): 234
```

group

The group command options manage groups.

Synopsis

group --help

group --list

group --info

group --status

group --edit
[--name *new_name*]
[--dnsserver *server*]
[--ntpserver *server*]
[--domainname *domain_name*]
[--autosupport {*yes|no*}]
[--support_tunnel {*yes|no*}]
[--smtp_server *smtp_server*]
[--smtp_port *smtp_port*]
[--smtp_auth {*yes|no*}]
[--smtp_username *user_name*]
[--smtp_encrypt_type {*none|starttls|ssl*}]
[--smtp_from_addr *email_addr*]
[--smtp_to_addr *email_addr*]
[--send_event_data {*yes|no*}]
[--alert_level {*info|notice|warning|critical*}]
[--proxyserver *server*]
[--proxyport *port*]
[--proxyuser *user_name*]
[--proxypasswd *password*]
[--default_vol_limit *percent*]
[--isns_enable {*yes|no*}]
[--isns_server *isns_server*]
[--isns_port *isns_port*]
[--snmp_trap_enabled {*yes|no*}]
[--snmp_trap_host *snmp_trap_hostname*]
[--snmp_trap_port *snmp_trap_port*]
[--snmp_get_enabled {*yes|no*}]
[--snmp_community *snmp_community*]
[--snmp_get_port *snmp_get_port*]
[--snmp_sys_location *snmp_sys_location*]
[--snmp_sys_contact *snmp_sys_contact*]
[--syslog_enabled {*yes|no*}]
[--syslog_server *syslog_server*]
[--syslog_port *syslog_port*]
[--encryption_cipher {*aes-256-xts|none*}]
[--encryption_scope {*group|volume*}]
[--encryption_mode {*available|secure*}]
[--inactivity_timeout *inactivity timeout*]
[--iscsi_target_name *iscsi_target_name*]
[--login_banner]
[--login_banner_after_auth {*yes|no*}]
[--force]
[--tlsvl_enabled {*yes|no*}]
[--cc_enabled {*yes|no*}]
[--alarms_enabled {*yes|no*}]

[--vss_validation_timeout *vss_validation_timeout*]
[--snap_ttl {*ttl*|*unlimited*}]
[--snap_ttl_unit {*hours*|*days*|*weeks*}]
[--iscsi_enabled {*yes*|*no*}]
[--fc_enabled {*yes*|*no*}]
[--tdz_enabled {*yes*|*no*}]
[--tdz_prefix *tdz_prefix*
[--analytics_gui {*yes*|*no*}]
[--default_iscsi_target_scope {*volume*|*group*}]

group --create_throttle
[--description *text*]
[--days *days*]
[--at *time*]
[--until *time*]
[--bandwidth *limit*]

group --delete_throttle *id*

group --edit_throttle *id*
[--description *text*]
[--days *days*]
[--at *time*]
[--until *time*]
[--bandwidth *limit*]

group --autosupport_validate

group --autosupport_initiate

group --merge_validate *group_name*
[--username *user_name*]
--password *password*--passphrase]
[--skip_secondary_mgmt_ip_validation]

group --merge *group_name*
[--username *user_name*]
--password *password*
[--passphrase]
[--allow_lower_limits]
[--skip_secondary_mgmt_ip_validation]

group --unset_http_proxy

group --list_limits

group --reset_login_banner

group --autoclean_unmanaged_snapshots {*yes*|*no*}
[--snap_ttl {*ttl*|*unlimited*}]
[--snap_ttl_unit {*hours*|*days*|*weeks*}]

group --show_login_banner

group --check_migrate

group --migrate

group --check_takeover

group --takeover

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List groups that share the same subnet with the current group.	Guest
--info	N/A	Provide detailed information about the current group.	Guest
--status	N/A	Display status of the group.	Guest

Option	Arguments and Suboptions	Description	Role
--edit	[--name *name*]	Modify the name of a group. You must include this option in a multi-array configuration. If you include this option in a single-array configuration, you modify the group name to the specified *name* value.	Power User
	[--dnsserver *server*]	Modify the IP address of the DNS server used. Repeat this option to specify more than one DNS server IP address.	
	[--ntpserver *server*]	Modify the hostname or IP address for the NTP server that is used by a group.	
	[--domainname *domain_name*]	Modify the DNS domain name that is used by a group.	
	[--autosupport {*yes*|*no*}]	Enable or disable automatically sending AutoSupport data to the Nimble Storage Support team.	
	[--support_tunnel {*yes*|*no*}]	Enable or disable Nimble Storage Support to establish a secure connection to the group to gather diagnostic data.	
	[smtp_server *smtp_server*]		

Option	Arguments and Suboptions	Description	Role
		Modify the hostname or IP address of an SMTP server that the group uses to send email alerts and event data to specified recipients.	
	[--smtp_port *smtp_port*]	Modify the SMTP port number to use. Default: 25. You must also include the --smtp_server option when you modify the port number.	
	[--smtp_auth {*yes*|*no*}]	Enable or disable SMTP authentication. Default: no (disabled). If you enable SMPT encryption, then you must also enable SMTP authentication.	
	[--smtp_username *user_name*]		

Option	Arguments and Suboptions	Description	Role
		Modify the SMTP authentication username and password. The *user_name* value must start with an alphabetic character and can be up to 64 alphanumeric characters. When prompted, interactively specify a password for the SMTP account. The *password* value can be blank or it can be up to 255 printable characters. You can repeat the prior *password* value.	
	[--smtp_encrypt_type { *none*|*starttls*|*ssl*}]	Modify the level of encryption for SMTP. Default: none. Requires using SMTP authentication when encryption is enabled by specifying a value of *starttls* (start transport layer security) or *ssl* (secure sockets layer). The *starttls* type can be an appropriate choice if you implemented cloud-based email. The *ssl* type can be an appropriate choice if you have a secure SMTP server installed on your network.	
	[--smtp_from_addr *email_addr*]		

Option	Arguments and Suboptions	Description	Role
		Modify the email address from which the group sends alerts and event data. It does not have to be a valid email address, but it must have a valid email address format. Include the group name for easy identification and filtering.	
	[--smtp_to_addr *email_addr*]	Modify one or more email addresses for alert and event recipients. Repeat this option to specify as many TO email addresses as needed.	
	[--send_event_data {*yes*|*no*}]	Enable or disable sending event data to the Nimble Storage Support team. Default: yes (enabled). When enabled, you do not need to provide contact details for Nimble Storage Support.	
	[--alert_level {*info*|*notice*|*warning*|*critical*}]	Specify the minimum level at which to send alert email notifications.	
	[--proxyserver *server*]		

Option	Arguments and Suboptions	Description	Role
		Modify the hostname or IP address of the HTTP proxy server. If no proxy server is specified, a direct HTTP connection is assumed. The array uses the HTTP protocol to download software updates and post event data.	
	[--proxyport *port*]	Modify the port number that the HTTP server uses.	
	[--proxyuser *user_name*]	Modify the username of the HTTP proxy server.	
	[--proxypasswd *password*]	Modify the password that the HTTP proxy server uses.	
	[--default_vol_limit *percent*]	Modify the percentage that is used as the default value for space limit when you create a volume for a group.	
	[--isns_enable {*yes*|*no*}]	Enable or disable the Internet Storage Name Service (iSNS) management services for the group. The iSNS settings do not apply to Fibre Channel arrays.	
	[--isns_server *isns_server*]	Specify the iSNS server that is used to publish iSCSI endpoints.	
	[--isns_port *isns_port*]		

Option	Arguments and Suboptions	Description	Role
		Modify the iSNS server port number.	
	[--snmp_trap_enabled {*yes*|*no*}]	Enable or disable SNMP traps.	
	[--snmp_trap_host *snmp_trap_hostname*]	Modify the hostname of the SNMP trap destination. Only one SNMP trap is currently available.	
	[--snmp_trap_port *snmp_trap_port*]	Modify the port number of the SNMP trap destination. The default destination port is 162.	
	[--snmp_get_enabled {*yes*|*no*}]	Enable or disable SNMP Gets.	
	[--snmp_community *snmp_community*]	Modify the name of the SNMP community. The default community is public. You can use alphanumeric characters and the - (hyphen), : (colon), and . (period) special characters. The name must start with an alphanumeric character and can be up to 64 characters.	
	[--snmp_get_port *snmp_get_port*]	Modify the SNMP responder port number. The default get port number is 161.	
	[--snmp_sys_location *snmp_sys_location*]		

Option	Arguments and Suboptions	Description	Role
		Modify the SNMP system location that is used to respond to SNMPv2-MIB::sysLocation.	
	[--snmp_sys_contact *snmp_sys_contact*]	Modify the SNMP system contact that is used to respond to SNMPv2-MIB::sysContact.	
	[--syslog_enabled {*yes*|*no*}]	Enable or disable sending syslog messages from the Nimble group to a remote syslog server. Default: no (syslog is disabled).	
	[--syslog_server *syslog_server*]	Modify the name of the server to which syslog messages will be forwarded.	
	[--syslog_port *syslog_port*]	Modify the remote syslog server port number. Default: 514. Valid values: 1 through 65535, inclusive.	
	[--encryption_cipher {*aes-256-xts*|*none*}]	Specify the default cipher to use during encryption. A value of "none" indicates no encryption. A value of "aes-256-xts" indicates the use of the AES-256-XTS cipher.	
	--encryption_scope {*group*|*volume*}]		

Option	Arguments and Suboptions	Description	Role
		Specify whether the default encryption_cipher setting can be overridden when creating a new volume, or is forced to be used. A value of "group" indicates that the value is applied across the entire group and cannot be changed. A value of "volume" means that the setting can be overridden at volume creation time.	
	[--encryption_mode {*available*|*secure*}]	Specify whether the master key passphrase must be entered on array startup ("secure" mode) to access encrypted volumes, or whether the array remembers the master key across a restart ("available" mode).	
	[--inactivity_timeout *inactivity timeout*]	Change the global user inactivity timeout (in minutes) that must pass without any user activity before a timeout occurs and users are logged out.	
	[--iscsi_target_name *iscsi_target_name*]		

Option	Arguments and Suboptions	Description	Role
		Changes the group-wide iSCSI target name. All Virtual Volumes and OpenstackV2 volumes are required to be offline. This is a disruptive operation.	
	[--alarms_enabled {*yes*|*no*}]	Enable/Disable alarms.	
	[--vss_validation_timeout {*vss_validation_timeout*}]	Changes the validation timeout for Microsoft VSS application synchronization (in minutes).	
	[--snap_ttl {*ttl*|*unlimited*}]	Time-to-live of an unmanaged snapshot from the time it is created. This options should be used with the --snap_ttl_unit option. The value that can be configured should be in the appropriate range, for example, 1 hour, 104 weeks.	
	[-snap_ttl_unit {*hours*|*days*|*weeks*}]	Unit of time used for the snapshot time to live option. This can be one of hours, days, or weeks.	
	[--iscsi_enabled {*yes*|*no*}]		

Option	Arguments and Suboptions	Description	Role
		Enable the iSCSI protocol for the group. Valid values are yes and no. This option cannot be used together with the --fc_enabled option. The change may cause some system processes to restart, which will impact I/O.	
	[--fc_enabled {*yes*|*no*}]	Enable the FC protocol for the group. Valid values are yes and no. This option cannot be used together with the --iscsi_enabled option. The change may cause some system processes to restart, which will impact I/O.	
	--tdz_enabled{*yes*|*no*}	Enable or disable Target Driven Zoning.	
	--tdz_prefix*tdz_prefix*	Target Driven Zoning (TDZ) prefix for peer zones created by TDZ.	
	[--gui_analytics {*yes*|*no*}]	Enable or disable usage analytic on the array.	
	[--default_iscsi_target_scope {*volume*|*group*}]		

Option	Arguments and Suboptions	Description	Role
		Export newly created volumes under iSCSI Group Target or iSCSI Volume Target. The valid options are volume and group. This setting is used only when creating a new volume without specifying the iSCSI default target scope for the volume.	
	[--login_banner]	Login banner message that is displayed during the user login process.	Administrator
	[--login_banner_after_auth {*yes*|*no*}]	Display login banner message before or after user authentication.	
	[--force]	The --force option confirms that no login banner exists. This option is mandatory to set an empty login banner.	
	[--tlsv1_enabled {*yes*|*no*}]	Enable or disable TSLv1.0 and TLSv1.1 in all SSL services. The change may cause some system processes to restart, but I/O will not be impacted.	
	[--cc_mode_enabled {*yes*|*no*}]		

Option	Arguments and Suboptions	Description	Role
		Enable or disable Common Criteria mode to control various configuration values. The change may case some system processes to restart, but I/O will not be impacted.	

Option	Arguments and Suboptions	Description	Role
--create_throttle	Specify at least one suboption.	Create a group-wide bandwidth throttle for replication. Throttles at the group level limit the total replication bandwidth across all downstream partners. The system automatically assigns an ID to the throttle schedule.	Power User
	[--description *text*]	Describe the bandwidth throttle. If the description includes spaces, enclose the description text in quotation marks.	
	[--days {*all* |*days*}]	Specify the days of the week for which the bandwidth throttle is effective. Separate the values with commas. For example: `Mon,Tue,Thu` Specify the value *all* to indicate all days of the week (the default).	
	[--at *time*]	Specify the time of day when the bandwidth throttle is activated.	
	[--until *time*]	Specify the time of day when the bandwidth throttle is deactivated.	
	[--bandwidth *limit*]		

Option	Arguments and Suboptions	Description	Role
		Specify the maximum bandwidth to use for replication, expressed in megabits per second (Mbps) or kilobits per second (Kbps). The default is megabits per second. Examples of valid arguments are 20Mbps or 400Kbps. Use 0 (zero) to completely throttle replication traffic. Use "unlimited" to turn off throttling.	
--delete_throttle	*id*	Delete the specified bandwidth throttle for replication. The ID is assigned by the system. Use group--info to view all throttles with their IDs.	Power User

Option	Arguments and Suboptions	Description	Role
--edit_throttle	*id*	Modify the specified bandwidth throttle for replication. Use group --info to determine the throttle ID. Specify at least one suboption.	Power User
	[--description *text*]	Modify the description of the bandwidth throttle. If the description includes spaces, enclose the description text in quotation marks.	
	[--days {*all* |*days*}]	Modify the days of the week for which the bandwidth throttle is effective. Separate the values with commas. For example: `Mon,Tue,Thu` Specify the value *all* to indicate all days of the week (the default).	
	[--at *time*]	Modify the time of day when the bandwidth throttle is activated.	
	[--until *time*]	Modify the time of day when the bandwidth throttle is deactivated.	
	[--bandwidth *limit*]		

Option	Arguments and Suboptions	Description	Role
		Modify the maximum bandwidth to use for replication, expressed in megabits (Mbps) or kilobits per second (Kbps). The default is megabits per second. Examples of valid arguments are 20Mbps or 400Kbps. Use 0 (zero) to completely throttle replication traffic. Use "unlimited" to turn off throttling.	
--autosupport_validate	N/A	Validate the AutoSupport configuration for a group.	Power User
--autosupport_initiate	N/A	Immediately send event information from the group to the Nimble Storage Support team. You can use this option in addition to configuring periodic AutoSupport messages by using the --edit --autosupport *yes* option.	Power User

Option	Arguments and Suboptions	Description	Role
--merge_validate	*group_name*	Validate the configuration of the specified group to merge it with the current group. This operation checks group objects for conflicts. If conflicts occur, then the information is listed so that you can resolve the conflicts before merging the groups.	Administrator
	[--username *user_name*]	Specify the administrative username associated with the group to validate before merging with this group. The default value is *admin*.	
	--password *password*	Specify the administrative password associated with the group to validate before merging with this group.	
	[--passphrase]	If encryption is enabled on the group that is the source of the merge being validated, specify the encryption passphrase for the source group.	
	[--skip_secondary_mgmt_ip_validation]	Use this option to skip the check whether the secondary management IP is missing.	

Option	Arguments and Suboptions	Description	Role
--merge	*group_name*	Merge the specified group into the current group. Use group --merge_validate to check for conflicts before merging two groups. After the groups are merged, all arrays from the merged group are placed in a new pool in the current group. Administration of the arrays and volumes on them must be performed in this group after the merge is complete.	Administrator
	[--username *user_name*]	Specify the administrative username associated with the group to merge with this group. The default value is *admin*.	
	--password *password*	Specify the group administrative password.	
	[--passphrase]	If encryption is enabled on the group that is the source of the merge, specify the encryption passphrase for the source group.	
	[--allow_lower_limits]	Allow group to have lower limits after the merge operation is complete.	

Option	Arguments and Suboptions	Description	Role
	--skip_secondary_mgmt_ip_validation	Skip the validation for the secondary management IP address.	
--unset_http_proxy	N/A	Clear the HTTP proxy server settings from the current group. The group will no longer use an HTTP proxy to connect to the Nimble Storage Support team.	Power User
--list_limits	N/A	List the current count, maximum limit, and warning threshold for several group criteria, as well as scope, where applicable.	Guest
--reset_login_banner	N/A	Resets the login banner to the factory default message.	Administrator

Option	Arguments and Suboptions	Description	Role
--autoclean_unmanaged_snapshots	{*yes*|*no*}	Enabled or disable the automatic cleaning of unmanaged snapshots.	Administrator
	[--snap_ttl {*ttl*|*unlimited*}]	Time-to-live of an unmanaged snapshot from the time it is created. This options should be used with the --snap_ttl_unit option. The value that can be configured should be in the appropriate range, for example, 1 hour, 104 weeks.	
	[--snap_ttl_limit {*hours*|*days*|*weeks*}]	Unit of time used for the snapshot time to live option. This can be one of hours, days, or weeks.	
--show_login_banner	N/A	Displays the login banner message.	Administrator
--check_migrate	N/A	Determine whether it is OK to migrate the group leader array to the backup group leader array. This command must be run from the array currently hosting the group Management Service (the group leader).	Power User

Option	Arguments and Suboptions	Description	Role
--migrate	N/A	Migrate data from the group leader array to the backup group leader array. This command must be run from the array currently hosting the group Management Service (the group leader). When this command completes, the backup group leader array will be the group leader array and will perform the management functions for the group.	Power User
--check_takeover	N/A	Determine whether it is OK to execute a takeover of the group leader array and promote the backup group leader array to group leader array. This command must be executed on the backup group leader. You must be logged on to the backup group leader array using at least a power user role.	Power User

Option	Arguments and Suboptions	Description	Role
--takeover	N/A	Takeover the role of group leader by the backup group leader. This command must be run from the Group Management Service backup array (the backup group leader). You must be logged on to the backup group leader array using at least a power user role. When this command completes, the backup group leader array will be the group leader array and will manage the group management services for the group.	Power User

Examples

This example lists the groups on a shared subnet.

```
Nimble OS $ group --list
------------------------------+---------------+--------------------------------
Name                          Mgmt IP Address  Version
------------------------------+---------------+--------------------------------
group-array7                  192.0.2.83       2.3.1.0-234720-opt
group-array8                  192.0.2.86       2.3.1.0-234720-opt
group-array3                  192.0.2.59       2.3.1.0-234619-opt
group-array5                  192.0.2.77       2.2.3.0-155034-opt
spark                         192.0.2.80       2.2.6.0-229590-opt
INFO: Only groups that share a subnet with this group (group-array2) are
discovered.
```

This example shows detailed information about the group that the array where you are logged in to the CLI is a member. The output is truncated to save space.

```
Nimble OS $ group --info
Group name: group-array2
Group Management IP: 192.0.2.51/255.255.255.0
Group Discovery  IP: 172.18.120.55/255.255.255.0
Group leader array: array2
Member array(s): array2
```

```
Unreachable member array(s):
Version: 2.3.1.0-234720-opt
NTP server: time.nimblestorage.com
Timezone: America/Los_Angeles
SMTP server:
SMTP port: 25
SMTP authentication: No
SMTP username:
SMTP password: (on file if set)
SMTP encryption type: NONE
Minimum alert email severity level: WARNING
Alert email recipient address(es):
Alert email from address:
Send event data to Nimble Storage Support: Yes
Send AutoSupport data to Nimble Storage Support: Yes
Allow Nimble Storage Support to establish secure connection to array: Yes
...
...
User inactivity timeout: 90 minute(s)
Encryption master key set: No
Encryption default settings scope: none
Encryption default cipher: none
Encryption security mode: none
Encryption activated: No
```

This example shows detailed information about the same group after encryption is configured. The output is truncated to save space.

```
Nimble OS $ group --info
Group name: group-array2
Group Management IP: 192.0.2.51/255.255.255.0
Group Discovery  IP: 172.18.120.55/255.255.255.0
Group leader array: array2
Member array(s): array2
Unreachable member array(s):
Version: 2.3.1.0-234720-opt
...
...
User inactivity timeout: 90 minute(s)
Encryption master key set: Yes
Encryption default settings scope: group
Encryption default cipher: aes-256-xts
Encryption security mode: secure
Encryption activated: Yes
```

This example modifies the group name and AutoSupport configuration.

```
Nimble OS $ group --edit --name greyhound2 --autosupport yes
```

This example configures secure SMTP for the group.

```
Nimble OS $ group --edit
--smtp_server mail.companyname.com --smtp_port 25
--smtp_auth yes --smtp_username smtp.user1
--smtp_encrypt_type ssl --smtp_from_addr group_name@companyname.com
--smtp_to_addr nimble.admin@companyname.com --send_event_data yes
Enter password:
Nimble OS $
```

This example creates a bandwidth throttle that limits replication bandwidth to 10 Mbps on Tuesdays from 6:00 AM until 7:00 AM.

```
Nimble OS $ group --create_throttle --description morning --days Tue
  --at 06:00AM --until 07:00AM --bandwidth 10
```

This example modifies the bandwidth throttle with the ID of 8 to permit 20 mbps during the permitted replication time.

```
Nimble OS $ group --edit_throttle 8 --bandwidth 20
```

This example deletes the bandwidth throttle for replication with the ID of 8.

```
Nimble OS $ group --delete_throttle 8
```

This example validates the AutoSupport configuration for a group with two arrays.

> **Note:** If the validation fails, an error message reports a potential reason. In some cases, an INFO message might suggest commands to use to diagnose the issue.

```
Nimble OS $ group --autosupport_validate
INFO: Successfully validated autosupport configuration on array array5
INFO: Successfully validated autosupport configuration on array array7
INFO: Successfully validated autosupport configuration.
```

This example sends AutoSupport information from the group to Nimble Storage Support.

```
Nimble OS $ group --autosupport_initiate
```

This example validates two groups before using the --merge option to merge the groups.

```
Nimble OS $ group --merge_validate group21 --user admin
--password dsfgH456
```

This example merges the group named *Seattle* into the current group.

```
Nimble OS $ group --merge Seattle --user admin
--password dsfgH456
```

This example clears the HTTP proxy server settings for the current group.

```
Nimble OS $ group --unset_http_proxy
```

This example lists limits for the group. The output is truncated to save space.

```
Nimble OS $ group --list_limits
---------------+---------------+----------------------+-------+-------+---------
Object          Scope Type      Scope Name             Count   Warning Maximum
---------------+---------------+----------------------+-------+-------+---------
array           group                                   1       4       4
array           pool                                    -       4       4
array           pool            default                 1       4       4
chapuser        group                                   0       960     1024
initiator       group                                   0       9500    10000
initiator       initiatorgrp                            -       240     256
initiatorgrp    group                                   1       960     1024
partner         group                                   0       8       10
partner         pool                                    -       8       10
perfpolicy      group                                   15      45      50
pool            group                                   1       4       4
protsched       group                                   6       950     1000
```

```
...
...
snapshot        group                                   0       36000   40000
snapshot        pool            default                 0       9000    10000
snapshot        vol                                     -       900     1000
snapcoll        volcoll                                 -       900     1000

NOTE: Snapshot counts for pool/group scope are for high retention volumes only.
```

halt

The halt command options halt an entire group, an array, or both controllers on a specified array. If you halt an array, it stops both controllers and turns off the array. To restart the array or controller, you have to press the Power button on the front of the array.

Synopsis

halt --help

halt
[--array *array_name*]
[--controller *{A | B|*]
[--group
[--force]
[--non_interactive]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--array	*array_name*	Halt the specified array. If you halt an array, it stops both controllers and turns off the array. To restart the array, press the Power button on the front of the array.	Power User
	[--controller *{A	B|*]	Halt the specified controller on the specified array. Valid controller names are A or B, which are case sensitive.. You can omit the --controller option if there is a single array in a group, an uninitialized array, or if entered on a member array in a group.	
	--non_interactive	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly taking the array or controller offline.	

Option	Arguments and Suboptions	Description	Role
--group		Halt the entire group when you run this command on the group leader. Both controllers of all arrays in a group are then halted and power is automatically turned off. If any array fails to respond to the halt request, then the group leader array remains operational. To restart the arrays in the group, press the Power button on the front of each array.	Power User
	--non_interactive	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly taking the group or controller offline.	
	[--force]	Halt the remaining arrays in a group when one or more arrays cannot be reached. For example, an array can be unreachable if it is already halted.	

Examples

This example halts an array named *NM127f*.

```
Nimble OS $ halt --array NM127f
WARNING: This operation will halt the array. Restarting the array will
require physically powering it back on. Type yes to continue, no to
cancel: yes
INFO: The controller shelf is being turned off. Therefore you can no longer
manage the group.
INFO: After the controller shelf has been turned off, you can turn off any
expansion shelves that are attached to the array.
INFO: Power on the expansion shelves first. Then power on the controller
shelf and log in to start managing the group again.
```

This example halts all controllers on all arrays in a group named *Acme_cluster*. You do not need to specify the group name. When you issue this command option on the group leader, it runs on the group that you are logged in to.

```
Nimble OS $ halt --group
WARNING: This operation will halt all arrays in the group.
Restarting the group will require physically powering on
all arrays in the group. Type yes to continue, no to cancel: yes
INFO: The controller shelf is being turned off. Therefore you can no longer
manage the group.
INFO: After the controller shelf has been turned off, you can turn off any
expansion shelves that are attached to the array.
INFO: Power on the expansion shelves first. Then power on the controller
shelf and log in to start managing the group again.
```

initiatorgrp

The initiatorgrp command options manage array access by grouping initiators that share a set of common characteristics that can be configured as part of your ACL to access specific volumes through group membership.

A given HPE Nimble Storage array supports access through the iSCSI or Fibre Channel (FC) protocol, and beginning with NimbleOS 5.1.x, both iSCSI and FC at the same time.

- An iSCSI initiator group is a set of initiators that are allowed to access a specific volume. After the initiator group is created, add iSCSI initiators to the group. All initiators in the group are granted access to the assigned volume(s).

  For example, to create an initiator group named *exchange* and assign an iSCSI initiator to it, you need to run two commands: one to create the initiator group and one to add the iSCSI initiators to it.

- A Fibre Channel (FC) initiator group is a set of initiators that are allowed to access specific volumes, or LUNs. The World Wide Port Names (WWPNs) for the host bus adapters (HBAs) on a host machine are used to create an initiator group. If a given WWPN is not in an initiator group that has an ACL to a volume (LUN), that host cannot access the volume.

Synopsis

initiatorgrp --help

initiatorgrp --list

initiatorgrp --info *initiatorgrp_name*

initiatorgrp --create *initiatorgrp_name*
[--description *text*]
[--access_protocol {*iscsi|fc*}
[--host_type *host_type*]

initiatorgrp --delete *initiatorgrp_name*

initiatorgrp --edit *initiatorgrp_name*
[--name *new_name*]
[--description *text*]
[--host_type *host_type*]

initiatorgrp --add_initiators *initiatorgrp_name*
[--label *label*]
[--initiator_name *iqn*]
[--ipaddr *ipaddr*]
[--initiator_alias *alias*]
[--wwpn *wwpn*]
[--force]

initiatorgrp --remove_initiator *initiatorgrp_name*
[--label *label*]
[--initiator_alias *alias*]
[--wwpn *wwpn*]

initiatorgrp --add_tdz_port *initiatorgrp_name*
[--interface_name *interface_name*]
[--array *{name|serial}*]

initiatorgrp --remove_tdz_port *initiatorgrp_name*
[--interface_name *interface_name*]
[--array *{name|serial}*]

initiatorgrp --add_subnets *initiatorgrp_name*
--label *subnet_label*

initiatorgrp --remove_subnet *initiatorgrp_name*

--label *subnet_label*
[--force]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all configured initiator groups.	Operator
--info	*initiatorgrp_name*	Show detailed information about the specified initiator group.	Operator
--create	*initiatorgrp_name*	Create a new iSCSI or Fibre Channel initiator group with the specified name.	Operator
	[--description *text*]	Add a description for the new initiator group. If there are spaces in the description, enclose the description text in quotation marks.	
	[--access_protocol {*iscsi*|*fc*}	Access protocol of initiator group. If this option is not specified, the access protocol of the supported group is used. If the array group supports multiple protocols, the access is set to Fibre Channel be default.	
	[--host_type *host_type*]	Set the host type for an initiator group. Available options are *auto*and *hpux*. The default is auto. This attribute will be applied to all the initiators in the initiator group. △ **CAUTION:** Initiators with different host OSes can only be kept in a group with the default (auto) host type attribute.	
--delete	*initiatorgrp_name*	Delete the specified initiator group. If any volumes include this initiator group in access control records, run initiatorgrp --remove_initiators to remove the group from those records first.	Operator

Option	Arguments and Suboptions	Description	Role
--edit	*initiatorgrp_name*	Modify attributes of the specified initiator group.	Operator
	[--name *new_name*]	Modify the name of the initiator group.	
	[--description *text*]	Modify the description of the initiator group. If there are spaces in the description, enclose the text in quotation marks.	
	[--host_type *host_type*]	Modify the host type of an initiator group. Available options are *auto* (the default) and *hpux*. This attribute will be applied to all the initiators in the initiator group. **CAUTION:** Initiators with different host OSes can only be kept in a group with the default (auto) host type attribute.	

Option	Arguments and Suboptions	Description	Role
--add_initiators	*initiatorgrp_name*	Add one or more members to the specified initiator group.	Operator
	[--label *label*]	An identifier of a minimum of two characters for an iSCSI initiator. The minimum allowed length is two characters. The label describes a set of attributes made up of the initiator name, IP address, and subnet combination (also called an access control list, or ACL).	
	[--initiator_name *iqn*]	Add an iSCSI initiator with this iSCSI Qualified Name (IQN) to the specified initiator group. Repeat this option to specify more than one initiator name. Each initiator name must have an associated IP address specified using the--ipaddr option. You can leave the IP address blank for the initiator if you prefer not to authenticate using both the name and IP address. The minimum allowed length of the IQN is two characters.	
	[--ipaddr *ipaddr*]	Add an iSCSI initiator with this IP address to the specified initiator group. Repeat this option to specify more than one initiator IP address, or specify "*" to match all IP addresses. Each initiator IP address must have an associated name specified using the --initiator_name option. Leave the name blank if you prefer not to authenticate using both name and IP address.	
	[--initiator_alias *alias*]	Add a Fibre Channel initiator with this alias of up to 32 characters to the specified initiator group. Repeat this option to specify more than one initiator alias. Each initiator alias must have an associated World Wide Port Name (WWPN) specified using the --wwpn option. You can choose to leave the WWPN for an initiator blank when using a previously saved initiator alias.	
	[--wwpn *wwpn*]		

Option	Arguments and Suboptions	Description	Role
		Add a Fibre Channel initiator with this WWPN to the specified initiator group. Repeat this option to specify more than one initiator WWPN. Each initiator WWPN can have an associated alias specified using the --initiator_alias option. Leave the alias blank if you prefer not to assign an initiator alias. Even if the alias is blank, if the specified WWPN and alias pair are used in another initiator group, the system will apply the alias when searching for, and finding, the same mapped values used elsewhere. The alias and WWPN have system-level scope.	
	[--force]	Update or remove conflicting Fibre Channel initiator aliases and forcibly adds the Fibre Channel initiator to the specified initiator group.	
--remove_initiator	*initiatorgrp_name*	Remove an initiator from the specified initiator group. Specify an iSCSI or Fibre Channel suboption as appropriate for the initiator group.	Operator
	[--label *label*]	Remove an iSCSI initiator with this unique label from the specified initiator group. The label describes a set of attributes made up of the initiator name, IP address, and subnet combination (also called an access control list, or ACL).	
	[--initiator_alias *alias*]	Remove a Fibre Channel initiator with this alias from the specified initiator group. If both initiator_alias and wwpn are provided, they must be synchronized. If they are not synchronized, the initiator will not be removed from the initiator group.	
	[--wwpn *wwpn*]	Remove a Fibre Channel initiator with this WWPN from the specified initiator group. If both --initiator_alias and --wwpn are provided, they must be synchronized. If they are not synchronized, the initiator will not be removed from the initiator group.	
--add_tdz_ports	*initiatorgrp_name*	Add a Target Driven Zone port to the specified initiator group.	
--add_tdz_ports	*initiatorgrp_name*	Add a Target Driven Zone port to the specified initiator group.	Operator
	--interface_name*interface_name*	Name of the Fibre Channel interface to be added (mandatory).	

Option	Arguments and Suboptions	Description	Role
	--array*{name|serial}*	Fibre Channel interface on the specified array.	
--remove_tdz_ports	*initiatorgrp_name*	Remove a Target Driven Zone port from the specified initiator group.	
	--interface_name*interface_name*	Name of the Fibre Channel interface to be removed from this initiatorgrp (mandatory).	
	--array*{name|serial}*	Fibre Channel interface on the specified array.	
--add_subnets	*initiatorgrp_name*	Add target subnets for the specified iSCSI initiator group to discover and access volumes.	Operator
	--label *subnet_label*	Add an iSCSI target subnet with this label to the specified initiator group. Use subnet --list to get a list of subnets. Repeat this option to specify more than one subnet.	
--remove_subnets	*initiatorgrp_name*	Remove a target subnet from the specified iSCSI initiator group.	Operator
	--label *subnet_label*	Remove an iSCSI target subnet with this label from the specified initiator group. Use initiatorgrp --list to see which initiator groups have associated subnets. Then use initiatorgrp --info *initiatorgrp_name* to get a list of potential subnets to remove.	
	[--force]	Forcibly remove a target subnet from the specified initiator group.	

Examples

This example shows a list of current initiator groups on an iSCSI HPE Nimble Storage array. A value of All in the Number of Subnets column means that no subnets are associated with the initiator group.

```
Nimble OS $ initiatorgrp --list
------------------------------+--------------------------+-----------------
Initiator Group Name           Number of Initiators       Number of Subnets
------------------------------+--------------------------+-----------------
initiatorgroup1                2                                 1
initiatorgroup2                7                               All
```

This example shows a list of current initiator groups on a Fibre Channel Nimble array.

```
Nimble OS $ initiatorgrp --list
------------------------------+--------------------------+-----------------
Initiator Group Name           Number of Initiators       Number of Subnets
------------------------------+--------------------------+-----------------
a11-a14                        2
a11-a14-5                      4
a11-a15                        2
```

```
a12-al6                         2
a12-al7                         2
C22Win5                         2
C22win4                         2
a11-al2-RHEL                    2
a11-al3-OL                      2
esxcluster                      6
```

This example shows detailed information about the iSCSI initiator group named *initiatorgroup1*.

```
Nimble OS $ initiatorgrp --info initiatorgroup1
Name: initiatorgroup1
Description:
Access Protocol: iscsi
Created: Apr 10 2018 14:56:45
Last configuration change: Apr 30 2018 14:00:03
Number of Subnets: 1
      Subnet Label: Subnet-1.141.0.0
Number of Initiators: 2
      Initiator Label: volumetwo-v38270
            Initiator Name: iqn.2012-11.com.nimblestorage:volumetwo-v38270
            Initiator IP Address: *
      Initiator Label: volumeone-v40dc27
            Initiator Name: iqn.2012.com.nimblestorage:volumeone-v40dc27
            Initiator IP Address: *
```

This example shows detailed information about the Fibre Channel initiator group named *esxcluster*.

```
Nimble OS $ initiatorgrp --info esxcluster
Name: esxcluster
Description:
Access Protocol: fc
Created: Jul 23 2018 17:48:18
Last configuration change: Jul 23 2018 17:55:07
Number of Initiators: 6
      Initiator: A11_hba2 (21:00:00:0e:1e:19:60:90)
      Initiator: A11_hba3 (21:00:00:0e:1e:19:60:91)
      Initiator: Esx2_hba1 (21:00:00:0e:1e:19:60:40)
      Initiator: Esx2_hba2 (21:00:00:0e:1e:19:60:41)
      Initiator: Esx1_hba1 (10:00:00:90:fa:53:a0:a0)
      Initiator: Esx1_hba2 (10:00:00:90:fa:53:a0:a1)
```

This example creates an iSCSI initiator group named *grouplevel2* and then verifies that the group exists.

```
Nimble OS $ initiatorgrp --create grouplevel2 --description
"group level two iSCSI initiators"

Nimble OS $ initiatorgrp --list
------------------------------+---------------------------+-----------------
Initiator Group Name          Number of Initiators       Number of Subnets
------------------------------+---------------------------+-----------------
grouplevel2                   0                                 All
initiatorgroup1               2                                   1
initiatorgroup2               7                                 All

Nimble OS $ initiatorgrp --info grouplevel2
Name: grouplevel2
Description: group level two iSCSI initiators
Access Protocol: iscsi
Created: Jul 10 2018 11:06:47
```

```
Last configuration change: Jul 10 2018 11:06:47
Number of Subnets: All
Number of Initiators: 0
```

This example shows that there are no access control records for the iSCSI initiator group named *grouplevel2* and then deletes the group.

```
Nimble OS $ initiatorgrp --list
------------------------------+--------------------------+-----------------
Initiator Group Name          | Number of Initiators     | Number of Subnets
------------------------------+--------------------------+-----------------
grouplevel2                   | 0                        | All
initiatorgroup1               | 2                        |   1
initiatorgroup2               | 7                        | All

Nimble OS $ initiatorgrp --delete grouplevel2
```

This example modifies the name and description of the iSCSI initiator group named *grouplevel2* and then shows the changes.

```
Nimble OS $ initiatorgrp --list
------------------------------+--------------------------+-----------------
Initiator Group Name          | Number of Initiators     | Number of Subnets
------------------------------+--------------------------+-----------------
grouplevel2                   | 0                        | All
initiatorgroup1               | 2                        |   1
initiatorgroup2               | 7                        | All


Nimble OS $ initiatorgrp --edit grouplevel2 --name group-level-2
                      --description "second-level group"

Nimble OS $ initiatorgrp --list
------------------------------+--------------------------+------------------
Initiator Group Name          | Number of Initiators     | Number of Subnets
------------------------------+--------------------------+------------------
group-level-2                 | 0                        | All
initiatorgroup1               | 2                        |   1
initiatorgroup2               | 7                        | All

Nimble OS $ initiatorgrp --info group-level-2
Name: group-level-2
Description: second-level group
Access Protocol: iscsi
Created: Jul 10 2018 11:06:47
Last configuration change: Jul 10 2018 11:47:07
Number of Subnets: All
Number of Initiators: 0
```

This example adds one initiator member to the iSCSI initiator group named *group-level-2*.

```
Nimble OS $ initiatorgrp --add_initiators group-level-2 --initiator_name
iqn.2012-11.com-company --ipaddr 10.12.130.211
```

This example removes an initiator from the iSCSI initiator group named *group-level-2*.

```
Nimble OS $ initiatorgrp group-level-2 --remove_initiator --label iqn2
```

This example adds one subnet to the iSCSI initiator group named *group-level-2*.

```
Nimble OS $ initiatorgrp --add_subnets group-level-2 --label data1
```

This example removes the *data1* subnet from the iSCSI initiator group named *group-level-2*.

```
Nimble OS $ initiatorgrp --remove_subnet group-level-2 --label data1
```

ip

The ip command options manage IP configuration settings.

The array management IP address is a floating address. It is assigned to a physical port by the system. This IP address is used to access the GUI and CLI on the array. The controller uses this IP address to have an active role. The discovery IP address is also a floating address. It is assigned to a physical port by the system. SCSI initiators use this IP address to discover targets. This array discovery IP address and the management IP address can share the same IP address or have unique IP addresses.

Each controller on the HPE Nimble Storage array has multiple port pairs that are used as data ports. You can configure these ports as part of the network configuration. Data ports are paired to ensure high availability (HA). The number of port pairs depends on your model.

Each controller is also assigned a static IP address to have direct access to a controller.

Synopsis

ip --help

ip --list
[--netconfig *netconfig_name*]
[--array *array_name*]

ip --info *ipaddr*
[--netconfig *netconfig_name*]
[--array *array_name*]

ip --add *ipaddr*
[--netconfig *netconfig_name*]
[--array *array_name*]
--type *{data|management|secondary|support}*
[--nic *nic_name*]
[--ctrlr *{A|B}*]

ip --edit *ipaddr*
[--netconfig *netconfig_name*]
[--array *array_name*]
[--type *{data|discovery|management|secondary|support}*]
[--nic *nic_name*]
[--newaddr *ipaddr*]
[--ctrlr *{A|B}*]
[--force_ip_update]

ip --delete *ipaddr*
[--netconfig *netconfig_name*]
[--array *array_name*]
[--type *{data|discovery|management|secondary|support}*]
[--ctrlr *{A|B}*]
[--force_ip_update]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List all configured data, management, and diagnostic (support) IP addresses.	Guest
	[--netconfig *netconfig_name*]	List the data, management, and diagnostic (support) IP addresses for the specified network configuration. Use netconfig --list to get the names of network configurations. Typical names are *{active | secondary | draft}*.	
	[--array *array_name*]	List the data, management, and diagnostic (support) IP addresses for the specified array. Use array --list to get the names of arrays.	
--info	*ipaddr*	Show detailed information about the specified IP address.	Guest
	[--netconfig *netconfig_name*]	Show detailed information about the specified IP address in the specified network configuration. Use netconfig --list to get the names of configurations.	
	[--array *array_name*]	Show detailed information about the specified IP address on the specified array. Use array --list to get the names of arrays.	

Option	Arguments and Suboptions	Description	Role
--add	*ipaddr*	Add an IP address to the network configuration for an array.	Power User
	[--netconfig *netconfig_name*]	Add the specified data, management, or support (diagnostic) IP address to the specified network configuration.	
	[--array *array_name*]	Add the specified data, management, or support (diagnostic) IP address to the specified array. If you specify the name of the backup group leader, this command adds the IP address for that array.	
	--type *{data|management|secondary|support}*	Specify the type of IP address being added. If you specify *secondary* and the group already has a backup group leader, this command will fail.	
	[nic *nic_name*]	Add a data IP address to the specified NIC. Values depend on the hardware configuration, and can include eth1 through eth6 for the CS-Series or eth1, eth2, tg1, and tg2 for the CS-Series "G" configuration (for CS200 and CS400 models).	
	[--ctrlr *{A|B}*]	Add a support IP address to the specified controller.	

Option	Arguments and Suboptions	Description	Role
--edit	*ipaddr*	Modify the specified IP address on the network configuration for an array.	Power User
	[--netconfig *netconfig_name*]	Modify the specified data, discovery, management, or support (diagnostic) IP address for the specified network configuration.	
	[--array *array_name*]	Modify the specified data, discovery, management, or support (diagnostic) IP address for the specified array. If you specify the name of the backup group leader, this command allows you to edit the IP address for that array	
	[--type *{data|discovery|management|secondary|support}*]	Modify the type for the specified IP address. Valid values are *data*, *discovery* (although this type has no effect on a Fibre Channel array), *management*, *support*, or *secondary*. If you specify *secondary* and the group already has a backup group leader, this command will fail.	
	[--nic *nic_name*]	Modify a data IP address for this NIC. Values depend on the hardware configuration, and can include eth1 through eth6 for the CS-Series or eth1, eth2, tg1, and tg2 for the CS-Series "G" configuration (for CS200 and CS400 models).	
	[--newaddr *ipaddr*]	Modify a previously configured IP address.	
	[--ctrlr *{A|B}*]	Modify a support IP address for the specified controller.	
	[--force_ip_update]	Forcibly modify a data or discovery IP address in an active network configuration. The command fails unless this option is used.	

Option	Arguments and Suboptions	Description	Role
--delete	*ipaddr*	Delete the specified IP address from a network configuration on an array.	Power User
	[--netconfig *netconfig_name*]	Delete the specified data, discovery, management, or support (diagnostic) IP address from the specified network configuration.	
	[--array *array_name*]	Delete the specified data, discovery, management, or support (diagnostic) IP address from the specified array. If you specify the name of the backup group leader, this command deletes the IP address for that array.	
	[--type *{data|discovery|management|secondary|support}*]	Delete an IP address configuration. Valid values are *data*, *discovery*, *management*, *secondary*, or *support*.	
	[--ctrlr *{A|B}*]	Delete a support IP address from the specified controller.	
	[--force_ip_update]	Forcibly delete a data or discovery IP address in an active network configuration. The command fails unless this option is used.	

Examples

This example shows information about the array IP configuration.

```
Nimble OS $ ip --list
---------------+--------+------+----------+------------------
IP Address      NIC      Status Type       Controller
---------------+--------+------+----------+------------------
192.0.2.209     eth1     up     management A
192.0.2.209     eth1     up     discovery  A
198.51.100.209  eth1     up     data       A
198.51.100.210  eth2     up     data       A
198.51.100.211  eth3     up     data       A
198.51.100.212  eth4     up     data       A
192.0.2.209     eth1     up     support    A
192.0.2.210     eth1     up     support    B
182.18.16.142   eth1     up     support    A
182.18.16.206   eth1     up     backup     B
```

This example shows detailed information about an IP address in the active network configuration.

```
Nimble OS $ ip --info 192.0.2.209 --netconfig active
IP address: 192.0.2.209
NIC: eth1
Link status: up
Array: c20-array2
Controller: B
Type: management, discovery
```

This example adds a support IP address to controller B.

```
Nimble OS $ ip --add 10.12.148.209 --type support --ctrlr B
```

This example modifies a data IP address and lists the change to verify the modification.

```
Nimble OS $ ip --edit 198.51.100.209 --type data --newaddr 198.51.100.109

Nimble OS $ ip --list
---------------+--------+------+----------+------------------
IP Address      NIC      Status Type       Controller
---------------+--------+------+----------+------------------
192.0.2.209     eth1     up     management A
192.0.2.209     eth1     up     discovery  A
198.51.100.109  eth1     up     data       A
198.51.100.210  eth2     up     data       A
198.51.100.211  eth3     up     data       A
198.51.100.212  eth4     up     data       A
192.0.2.209     eth1     up     support    A
192.0.2.210     eth1     up     support    B
```

This example deletes the support IP address for controller B and then lists the IP addresses to verify the deletion.

```
Nimble OS $ ip --delete 192.0.2.210 --type support --ctrlr B

ip --list
---------------+--------+------+----------+------------------
IP Address      NIC      Status Type       Controller
---------------+--------+------+----------+------------------
192.0.2.209     eth1     up     management A
192.0.2.209     eth1     up     discovery  A
198.51.100.109  eth1     up     data       A
198.51.100.210  eth2     up     data       A
198.51.100.211  eth3     up     data       A
198.51.100.212  eth4     up     data       A
192.0.2.209     eth1     up     support    A
```

keymanager

The keymanager command is used to manage external key managers for encryption keys. Encrypted keys can be created and then stored on a secure server with a third-party key management provider. Managing keys off the array makes it more secure, and less vulnerable to security breaches.

External key managers provide a mechanism to securely store and retrieve sensitive information. Encryption keys from volumes are stored in an internal database. Adding an external key manager allows this information to be stored separately from the array.

Synopsis

keymanager --help

keymanager --list

keymanager --info *keymgr_name*

keymanager --edit *keymgr_name*
[--hostname *hostname*]
[--port *port*]

[--protocol {*KMIP1_0 | KMIP1_1 | KMIP1_2 | KMIP1_3*}]
[--username *user*]
[--password *pass*]

keymanager --add *keymgr_name*
--hostname *hostname* --port *port*
--protocol {*KMIP1_0 | KMIP1_1 | KMIP1_2 | KMIP1_3*}
[--username *user*]
[--password *pass*]

keymanager --migrate *keymgr_name*

keymanager --remove *keymgr_name*

Table

Option	Arguments and Suboptions	Description	Role
--help	N/A	Display command options and suboptions.	Operator
--list	N/A	List external key managers configured for the group.	Operator
--info	*keymgr_name*	Display information about a specific external keymanager.	Operator
--edit	*keymgr_name*	Modify the settings for the specified external key manager.	Administrator
	[--hostname *hostname*]	Modify the host name or IP address of the of the external key manager.	
	[--port *port*]	Modify the port number of the external key manager.	
	[--protocol {*KMIP1_0	KMIP1_1	KMIP1_2	KMIP1_3*}]	Modify the KMIP protocol to be used for communication with the external key manager	
	[--username *user*]	Modify the username for authentication with the external key manager.	
	[--password *pass*]	Modify the password for authentication with the external key manager	

Option	Arguments and Suboptions	Description	Role
--add	*keymgr_name*	Add an external key manager for the group.	Administrator
	--hostname *hostname*	The host name or IP address of the external key manager server. This is a mandatory argument for adding a keymanager.	
	--port *port*	Port number for the external key manager. This is a mandatory argument for adding a keymanager.	
	--protocol {*KMIP1_0 | KMIP1_1 | KMIP1_2 | KMIP1_3*}	KMIP protocol to be used for communication with the external key manager. This is a mandatory argument argument for adding a keymanager.	
	[--username *user*]	Username for authentication with the external key manager	
	[--password *pass*]	Password for authentication with external key manager.	
--migrate	--keymanager *keymgr_name*	Migrate encryption keys from an active external key manager to an inactive external key manager.	
--remove	--keymanager *keymgr_name*	Remove an external key manager key manager configuration. Encryption keys will be migrated to the internal key manager on the array.	

migration

The migration command options provide information about data migration activity.

Synopsis

migration --help

migration --list
[--array_unassign]
[--pool_rebal]
[--vf_reloc]

migration --info *{array|pool|volume}*
[--array_unassign]
[--pool_rebal]
[--vf_reloc]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--list	{--array_unassign|--pool_reball|--vf_reloc}	List data-migration activity that is related to unassigning arrays from pools, rebalancing array usage in pools, or relocating a volume family to a different pool. Specify one of the suboptions each time you run migration --list.	Guest
--info	{*array_name* --array_unassign|*pool_name* --pool_reball|*volume_name* --vf_reloc}	Show information about data-migration activity that is related to the specified array being unassigned from its pool, the specified pool being rebalanced, or the specified volume family being relocated to a different pool. A volume family is a collection of volumes with a single root volume. A root volume is a volume that has no parent. Specify one of the argument/suboption pairs each time you run migration --info.	Guest

Examples

This example shows that no migration of array unassignment is currently active.

```
Nimble OS $ migration --list --array_unassign
None.
```

This example shows information about unassigning an array from its pool.

```
Nimble OS $ migration --info c20-array2 --array_unassign
Array Unassign:
--------------------+----------------+----------------------
Array               Data             Estimated
                    Remaining (MB)   Completion
--------------------+----------------+----------------------
c20-array2                      932   2 min
```

netconfig

The netconfig command options manage a network configuration.

Synopsis

netconfig --help

netconfig --list

netconfig --info {*active|backup|draft*}

netconfig --create_draft_from {*active|backup*}
[--force]

netconfig --validate {*active* | *backup* | *draft*}
[--force_repl]
[--force_initiator_groups]

netconfig --activate {*draft|backup*}
[--force_repl]
[--force_ip_update]
[--force_initiator_groups]

netconfig --edit {*active|backup|draft*}
[--iscsi_automatic_connection_method {*yes|no*}]
[--iscsi_connection_rebalancing {*yes|no*}]
[--force_ip_update]

netconfig --delete {*backup|draft*}

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List network configurations.	Guest
--info	{*active|backup|draft* }	Provide detailed information about the specified network configuration.	Guest
--create_draft_from	{*active|backup*}	Create a draft network configuration from the specified active or backup network configuration.	Power User
	[--force]	Forcibly create a draft network configuration if a draft already exists.	
--validate	{*active|backup|draft*}	Validate the specified network configuration.	Power User
	[--force_repl]	Ignore inconsistent network configurations for replication partners when validating the specified network configuration.	
	[--force_initiator_groups]	Ignore initiators assigned to subnets when validating the specified network configuration.	
--activate	{*draft|backup*}	Activate the specified network configuration and makes it the newly *active* network configuration.	Power User
	[--force_repl]	Ignore inconsistent network configurations for replication partners when activating the specified network configuration.	
	[--force_ip_update]	Allow activating the specified network configuration and updating the discovery or data IP addresses for the active configuration.	
	[--force_initiator_groups]	Ignore initiators assigned to subnets when activating the specified network configuration.	

Option	Arguments and Suboptions	Description	Role
--edit	{*active*|*backup*|*draft*}	Modify the specified network configuration. Specify at least one suboption.	Power User
	[--iscsi_automatic_connection_method {*yes*|*no*}]	Redirect connections from a specified iSCSI target IP to the best data IP that is based on connection counts. This option has no effect on a Fibre Channel array.	
	[--iscsi_connection_rebalancing{ *yes*|*no*}]	Rebalance iSCSI connections by periodically breaking existing connections that are out of balance to allow the host to reconnect to another data IP address. This option has no effect on a Fibre Channel array.	
	[--force_ip_update]	Allow updating the discovery or data IP addresses for the *active* network configuration.	
--delete	{*backup*|*draft*}	Delete the specified network configuration. You cannot delete the active network configuration.	Power User

Examples

This example lists the network configurations for an array.

```
Nimble OS $ netconfig --list
--------+--------------------+--------------------+--------------------
Name     Last Modified         Active Since          Last Active
--------+--------------------+--------------------+--------------------
backup   Jun 27 2014 11:22:45  N/A                   Jun 27 2014 11:22:55
active   Jun 27 2014 11:22:55  Jun 27 2014 11:22:45  N/A
```

This example shows information about the active network configuration on a Fibre Channel array.

```
Nimble OS $ netconfig --info active
Group Management IP: 192.0.2.19
Group leader array: array7
Member array(s): array7
ISCSI Automatic connection method: Yes
ISCSI Connection rebalancing     : Yes

Routes:
---------------+---------------+---------------
Destination     Netmask         Gateway
---------------+---------------+---------------
0.0.0.0         0.0.0.0         203.0.113.1

Subnets:
--------------------+-----------------+-------+--------------+----+-----
Label                Network           Type    Discovery IP   VLAN MTU
--------------------+-----------------+-------+--------------+----+-----
data1                192.0.2.0/23      Data    192.0.2.61     0    9000
mgmt-only            192.0.2.0/21      Mgmt    192.0.2.19     0    1500
Subnet-1.101.0.0     1.101.0.0/16      Data    N/A            2101 1500
Subnet-1.102.0.0     1.102.0.0/16      Data    N/A            2102 1500
```

```
Subnet-1.103.0.0      1.103.0.0/16      Data    N/A            2103 1500
Subnet-1.104.0.0      1.104.0.0/16      Data    N/A            2104 1500
Subnet-1.105.0.0      1.105.0.0/16      Data    N/A            2105 1500
Subnet-1.106.0.0      1.106.0.0/16      Data    N/A            2106 1500
Subnet-1.107.0.0      1.107.0.0/16      Data    N/A            2107 1500
Subnet-1.108.0.0      1.108.0.0/16      Data    N/A            2108 1500
Subnet-1.109.0.0      1.109.0.0/16      Data    N/A            2109 1500
Subnet-1.110.0.0      1.110.0.0/16      Data    N/A            2110 1500

Array Network Configuration: array7
Controller A IP: 192.0.2.11
Controller B IP: 192.0.2.18
---------+------------------------+---------------+------
NIC       Subnet Label             Data IP Address Tagged
---------+------------------------+---------------+------
eth1      mgmt-only                N/A             No
eth2                               N/A             No
eth3      data1                    192.0.2.62      No
eth3.2101 Subnet-1.101.0.0         1.101.3.35      Yes
eth3.2102 Subnet-1.102.0.0         1.102.3.35      Yes
eth3.2103 Subnet-1.103.0.0         1.103.3.35      Yes
eth3.2104 Subnet-1.104.0.0         1.104.3.35      Yes
eth3.2105 Subnet-1.105.0.0         1.105.3.35      Yes
eth3.2106 Subnet-1.106.0.0         1.106.3.35      Yes
eth3.2107 Subnet-1.107.0.0         1.107.3.35      Yes
eth3.2108 Subnet-1.108.0.0         1.108.3.35      Yes
eth3.2109 Subnet-1.109.0.0         1.109.3.35      Yes
eth3.2110 Subnet-1.110.0.0         1.110.3.35      Yes
eth4      data1                    192.0.2.91      No
eth4.2101 Subnet-1.101.0.0         1.101.194.35    Yes
eth4.2102 Subnet-1.102.0.0         1.102.194.35    Yes
eth4.2103 Subnet-1.103.0.0         1.103.194.35    Yes
eth4.2104 Subnet-1.104.0.0         1.104.194.35    Yes
eth4.2105 Subnet-1.105.0.0         1.105.194.35    Yes
eth4.2106 Subnet-1.106.0.0         1.106.194.35    Yes
eth4.2107 Subnet-1.107.0.0         1.107.194.35    Yes
eth4.2108 Subnet-1.108.0.0         1.108.194.35    Yes
eth4.2109 Subnet-1.109.0.0         1.109.194.35    Yes
eth4.2110 Subnet-1.110.0.0         1.110.194.35    Yes
```

This example creates a draft from the active network configuration and then lists the results.

```
Nimble OS $ netconfig --create_draft_from active

Nimble OS $ netconfig --list
----------+---------------------+---------------------+---------------------
Name       Last Modified          Active Since           Last Active
----------+---------------------+---------------------+---------------------
backup     Jun 27 2014 11:22:45   N/A                    Jun 27 2014 11:22:55
active     Jun 27 2014 11:22:55   Jun 27 2014 11:22:45   N/A
draft      Jul 17 2014 15:14:34   N/A                    N/A
```

This example validates a *draft* network configuration.

```
Nimble OS $ netconfig --validate draft
INFO: Configuration is valid.
```

If the validation fails, an INFO message identifies what issues to resolve. For example, on an iSCSI array:

```
Nimble OS $ netconfig --validate backup
ERROR: Failed to validate netconfig. Invalid argument.
INFO: Must configure at least one iSCSI and at least one group enabled subnet.
```

This example activates the *draft* network configuration.

```
Nimble OS $ netconfig --activate draft
```

This example edits the *draft* network configuration on an iSCSI array to enable iSCSI connection rebalancing.

```
Nimble OS $ netconfig --edit draft --iscsi_connection_rebalancing yes
```

This example deletes the *draft* network configuration.

> **Note:** If you delete the *backup* network configuration, you have to modify the *active* configuration and activate the modified version.

```
Nimble OS $ netconfig --delete draft
```

nic

The nic command options manage Network Interface Cards (NICs) on a Nimble array.

NICs are the hardware that manages any physical ports on an array. You can manage settings, such as frame size, on the NIC.

> **Note:** Do not enable jumbo frames on switches unless jumbo frames are also configured on the NICs.

Synopsis

nic --help

nic --list
[--netconfig *name*]
[--array *{name|serial}*]

nic --info *name*
[--netconfig *name*]
[--array *{name|serial}*]

nic --assign *name*
[--netconfig *name*]
[--array *{name|serial}*]
[--subnet *subnet-label*]
[--tagged *{yes|no}*]
[--data_ip *addr*]

nic --unassign *name*
[--netconfig *name*]
[--array *{name|serial}*]
[--subnet *subnet-label*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List all NICs, or the NICs associated with the specified network configuration or the NICs on the specified array.	Guest
	[--netconfig *name*]	List the NICs for a specified network configuration.	
	[--array *{name | number}*]	List the NICs on the specified array.	
--info	*name*	Show detailed information about the specified NIC.	Guest
	[--netconfig *name*]	Show detailed information about the specified NIC on the specified network configuration.	
	[--array *{name | number}*]	Show detailed information about the specified NIC on the specified array.	
--assign	*name*	Assign a subnet on the specified NIC.	Power User
	[--netconfig *name*]	Assign a subnet on the specified NIC on the specified network configuration.	
	[--array *{name | number}*]	Assign a subnet on the specified NIC on the specified array.	
	[--subnet *subnet-label*]	Specify the subnet to assign. Run subnet --list to get currently unassigned values for *subnet-label*.	
	[--tagged *{yes|no}*]	Specify whether the NIC uses a tagged assignment on the specified subnet. A tagged assignment means that traffic for the subnet on this NIC will carry the 802.1Q frame with the VLAN ID for the subnet. A NIC can have at most one untagged assignment and zero or many tagged assignments.	
	[--data_ip *addr*]	Specify the data IP address.	
--unassign	*name*	Unassign a subnet from the specified NIC.	Power User
	[--netconfig *name*]	Unassign a subnet from the specified NIC on the specified network configuration.	
	[--subnet *subnet-label*]	Unassign a subnet from the specified NIC on the specified array.	

Examples

This example lists the NICs on an iSCSI array. The output is truncated to avoid line wrapping.

```
nic --list
-------+-------+-----+----------------+----------------+-----------------
Name    Link    Link  MAC               Subnet           IP Addresses
        Status  Speed
        A/B
-------+-------+-----+----------------+----------------+-----------------
eth1    up/up      1G 00:25:90:75:18:90 mgmt-data        192.0.2.51,...
eth2    up/up      1G 00:25:90:75:18:91 mgmt-data
eth3    up/up      1G 00:E0:ED:18:25:06 data1            198.51.100.56
eth4    up/up      1G 00:E0:ED:18:25:07 data1            198.51.100.57
eth5    up/up      1G 00:E0:ED:1C:FF:66 data1            198.51.100.55,...
eth6    up/up      1G 00:E0:ED:1C:FF:67 data1            198.51.100.59
```

This example lists the NICs on a Fibre Channel array. The output is truncated to avoid line wrapping and to save space.

```
nic --list
---------+-------+-----+----------------+----------------+-------------------
Name      Link    Link  MAC               Subnet           IP Addresses
          Status  Speed
          A/B
---------+-------+-----+----------------+----------------+-------------------
eth1      up/up       1G 00:25:90:8C:F2:3A mgmt-only        192.0.2.18,...
eth2      up/up       1G 00:25:90:8C:F2:3B N/A
eth3      up/up      10G 00:25:90:E0:A5:71 data1            192.0.2.61,...
eth3.2101 up/up      10G 00:25:90:E0:A5:71 Subnet-1.101.0.0 1.101.3.35
eth3.2102 up/up      10G 00:25:90:E0:A5:71 Subnet-1.102.0.0 1.102.3.35
eth3.2103 up/up      10G 00:25:90:E0:A5:71 Subnet-1.103.0.0 1.103.3.35
eth3.2104 up/up      10G 00:25:90:E0:A5:71 Subnet-1.104.0.0 1.104.3.35
eth3.2105 up/up      10G 00:25:90:E0:A5:71 Subnet-1.105.0.0 1.105.3.35
eth3.2106 up/up      10G 00:25:90:E0:A5:71 Subnet-1.106.0.0 1.106.3.35
eth3.2107 up/up      10G 00:25:90:E0:A5:71 Subnet-1.107.0.0 1.107.3.35
eth3.2108 up/up      10G 00:25:90:E0:A5:71 Subnet-1.108.0.0 1.108.3.35
eth3.2109 up/up      10G 00:25:90:E0:A5:71 Subnet-1.109.0.0 1.109.3.35
eth3.2110 up/up      10G 00:25:90:E0:A5:71 Subnet-1.110.0.0 1.110.3.35
eth4      up/up      10G 00:25:90:E0:A5:70 data1            10.18.122.91
eth4.2101 up/up      10G 00:25:90:E0:A5:70 Subnet-1.101.0.0 1.101.194.35
eth4.2102 up/up      10G 00:25:90:E0:A5:70 Subnet-1.102.0.0 1.102.194.35
```

This example shows detailed information for the *eth2* NIC on an iSCSI array.

```
Nimble OS $ nic --info eth2
Name: eth2
Location on Controller A: 0,1
Location on Controller B: 0,1
Present on Controller A: Yes
Present on Controller B: Yes
Link Status on Controller A: up
Link Status on Controller B: up
Max Supported Link Speed: 1G
Interface Type: TP
Negotiated Link Speed on Controller A: 1G
Negotiated Link Speed on Controller B: 1G
MAC Address on Controller A: 00:25:90:75:18:9D
MAC Address on Controller B: 00:25:90:75:18:91
Subnet: mgmt-data (192.0.2.0/255.255.255.0)
Type: Mgmt
VLAN Id: 0
MTU: 1500
IP Address:
```

This example shows detailed information for the *eth3.2101* NIC on a Fibre Channel array.

```
Nimble OS $ nic --info eth3.2101
Name: eth3.2101
Location on Controller A: 1,0
Location on Controller B: 1,0
Present on Controller A: Yes
Present on Controller B: Yes
Link Status on Controller A: up
Link Status on Controller B: up
Max Supported Link Speed: 10G
Interface Type: TP
Negotiated Link Speed on Controller A: 10G
Negotiated Link Speed on Controller B: 10G
MAC Address on Controller A: 00:25:90:CA:D5:29
MAC Address on Controller B: 00:25:90:E0:A5:71
Subnet: Subnet-1.101.0.0 (1.101.0.0/255.255.0.0)
Type: Data
VLAN Id: 2101
MTU: 1500
IP Address: 1.101.3.35
```

partner

The partner command options manage replication partners and per-partner bandwidth throttles.

When you create a replication partner, you enable one Nimble array to replicate its data to another. The two arrays must be able to communicate over a network by using ports 4213 and 4214. Replication partners have the same name as the remote array.

Replication partners can be reciprocal, upstream (the source of replicas), or downstream (the receiver of replicas) partners. The replication partners must be configured on both arrays to successfully replicate.

Synopsis

partner --help

partner --list

partner --info *partner_name*

partner --create *partner_name*
--hostname *{ipaddr|hostname}*
[--control_port *port*]
[--data_port *port*]
[--description *text*]
[--secret *shared_secret*]
[--pool *pool_name*]
[--folder *folder_name*--match_folder *{yes|no}*]
[]
[--subnet *label*]

partner --edit *partner_name*
[--hostname *{ipaddr|hostname}*]
[--control_port *port*]]
[--data_port *port*]
[--description *text*]
[--secret *shared_secret*]

[--pool *pool_name*]
[--folder *folder_name*]
[--match_folder *{yes|no}*]
[--name *new_name*]
[--subnet *label*]

partner --delete *partner_name*

partner --create_throttle *partner_name*
[--description *text*]
[--days *days*]
[--at *time*]
[--until *time*]
[--bandwidth *limit*]

partner --edit_throttle *partner_name*
[--description *text*]
[--id *number*]
[--days *days*]
[--at *time*]
[--until *time*]
[--bandwidth *limit*]

partner --delete_throttle *partner_name*
--id *number*

partner --pause *partner_name*

partner --resume *partner_name*

partner --test *partner_name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all configured replication partners for the array.	Guest
--info	*partner_name*	Provide detailed information about the specified replication partner.	Guest

Option	Arguments and Suboptions	Description	Role
--create	*partner_name*	Create a replication partner. The partner name must be the array name. After creating a replication partner on one array, create a reciprocal, upstream, or downstream partner on the array configured as the replication partner. **Note:** With synchronous replication, a user cannot manually create pool partners. Pool partners are created based on the number of pools in the group.	Power User
	--hostname *{ipaddr|hostname}*	Specify the IP address or hostname of the management IP address for the replication partner.	
	[--control_port *port*]	Specify the port number of the control interface. If you do not specify this option, the default port 4213 is used.	
	[--data_port *port*]	Specify the port number of the data interface. If you do not specify this option, the default port 4214 is used.	
	[--description *text*]	Specify a description of the replication partner. If there are spaces in the description, enclose the text in quotation marks.	
	--secret *shared_secret*	Specify the shared secret that ensures secure communication between the replication partners. Both replication partners must use the same shared secret.	
	[--pool *pool_name*]	Specify the storage pool in which to create volumes. Any volumes replicated from this partner are created in the specified pool. Replica volumes created as clones ignore this parameter and are always created in the same pool as their parent volume. If you do not specify this option, the default storage pool is used.	
	[--folder *folder_name*]	Specify the name of the folder to contain the replica volumes. If the folder option is not specified, the replica volumes will be contained directly by the pool.	
	[--match_folder *{yes|no}*	Specify whether to match the upstream volume's folder on the downstream.	
	--subnet *label*	Specify the local subnet to use for communication with this partner. For example, 10.17.0.0/255.255.0.0.	

Option	Arguments and Suboptions	Description	Role
--edit	*partner_name*	Modify the configuration of the specified replication partner. Specify at least one suboption. **Note:** For synchronous replication, the only attributes of the partner command that you can edit are the description and the secret.	Power User
	[--hostname *{ipaddr|hostname}*]	Modify the IP address or hostname of the management IP address for the replication partner.	
	[--control_port *port*]	Modify the port number of the control interface. The default port number is 4213.	
	[--data_port *port*]	Modify the port number of the data interface. The default port number is 4214.	
	[--description *text*]	Modify the description of the replication partner.	
	[--secret *shared_secret*]	Modify the shared secret that ensures secure communication between the replication partners. Both replication partners must use the same shared secret.	
	[--pool *pool_name*]	Modify the storage pool in which to create volumes. Any volumes replicated from this partner are created in the specified pool. Replica volumes created as clones ignore this parameter and are always created in the same pool as their parent volume. Modifying the *pool_name* does not affect existing replica volumes.	
	[--folder *folder_name*]	Specify the name of the folder to contain the replica volumes. If the folder option is not specified, the replica volumes will be contained directly by the pool.	
	[--match_folder *{yes|no}*]	Specify whether to match the upstream volume's folder on the downstream.	
	[--new_name *partner_name*]	Modify the name of the replication partner. The name must match the group name of the replication partner	
	[--subnet *label*]	Modify the local subnet to use for communication with this partner. For example, 10.17.0.0/255.255.0.0.	

Option	Arguments and Suboptions	Description	Role
--delete	*partner_name*	Delete the specified replication partner. If a volume collection is being replicated to a replication partner, then subsequent replication based on that collection fails. **Note:** With synchronous replication, a user cannot manually delete pool partners. Pool partners are deleted based on the number of pools in the group.	Power User
--create_throttle	*partner_name*	Create a per-partner bandwidth throttle. You can create per-partner replication throttles only if there is no array-wide bandwidth throttle. You can create multiple per-partner throttles, but the throttles cannot overlap.	Operator
	[--description *text*]	Describes the throttle. The system automatically assigns an ID to the throttle schedule. If the description includes spaces, enclose the description text in quotation marks.	
	[--days *days*]	Specifies the days of the week on which the throttle is effective, separated by commas. For example, `Tue,Thur` You can specify the value *all* to indicate all days of the week, which is the default if you do not specify this option.	
	--at *time*	Specify the time of day at which the throttle is activated.	
	--until *time*	Specify the time of day at which the throttle is deactivated. This value must be later than the --at *time* value.	
	--bandwidth *limit*	Specify the maximum bandwidth to use for replication in megabits per second. Use 0 (zero) to stop replication traffic, or unlimited to turn off throttling.	

Option	Arguments and Suboptions	Description	Role
--edit_throttle	*partner_name*	Modify the specified per-partner bandwidth throttle.	Operator
	[--description *text*]	Modify the description of the throttle.	
	--id *number*	Specify the ID of the throttle to modify. Use the --info option to determine the throttle ID.	
	[--days *days*]	Modify the days of the week on which the throttle is effective, separated by commas.	
	--at *time*	Modify the time of day at which the throttle is activated.	
	--until *time*	Modify the time of day at which the throttle is deactivated.	
	--bandwidth *limit*	Modify the maximum bandwidth to use for replication, specified in megabits per second. Use 0 (zero) to stop replication traffic, or unlimited to turn off throttling.	
--delete_throttle	*partner_name*	Delete the specified per-partner bandwidth throttle.	Operator
	--id *number*	Specify the ID of the bandwidth throttle to delete. The ID is assigned by the system. Use partner --info to view all throttles and get the *number* you need.	
--pause	*partner_name*	Pause replication to or from the specified replication partner. In-progress replications are stopped until resumed. This state is stored persistently and remains in force across system restarts.	Operator
--resume	*partner_name*	Resume the paused replication to or from the specified replication partner.	Operator
--test	*partner_name*	Test the network connection to the specified replication partner. Both partners must be configured for the test to succeed.	Operator

Examples

This example lists configured replication partners.

```
Nimble OS $ partner --list
--------------------+------------------+---------------+----------
Partner             Hostname           Status          Throttled
                                                       Bandwidth
--------------------+------------------+---------------+----------
array7              array7             disconnected    unlimited
array1              192.0.2.93         connected       unlimited
array2              192.0.2.57         connected       unlimited
```

This example shows details about a synchronous replication partner named *array1*.

```
Nimble OS $ partner --info array1
Name: array1
Description: pool partner
Partner type: pool
Hostname: N/A
Control port: N/A
Data port: N/A
Pool: default
Folder:
Match Folder
Paused: No
Connected: Yes
Connection error: N/A
Configuration synchronized: N/A
Synchronization error:N/A
Throttled bandwidth: N/A
Throttles: N/A
Subnet type: N/A
Subnet: N/A
Array serial: N/A
Replication version: N/A
Created: Aug 4 2018 23:41:26
Last configuration change: April  5 2018 11:53:31
```

This example creates a replication partner with a secret of *nimblereplica* and a description of *San Jose array*. Because values for the control and data ports are not provided, the defaults of 4213 and 4214 are used.

```
Nimble OS $ partner --create array1 --description "San Jose array"
--secret nimblereplica --subnet 1.101.0.0/255.255.0.0
```

This example modifies the shared secret for a replication partner. In this case, the shared secret has to be changed on the replication partner, too.

```
Nimble OS $ partner --edit array1 --secret mycoreplication
```

This example deletes the replication partner named *array7*.

```
Nimble OS $ partner --delete array7
```

This example creates a per-partner bandwidth throttle for the replication partner *array2* that allows replication traffic Monday through Wednesday from 2 PM until 5 PM.

```
Nimble OS $ partner --create_throttle array2 --description "afternoon limit"
--days Mon,Tue,Wed --at 14:00 --until 17:00 --bandwidth 10
```

This example modifies a per-partner bandwidth throttle with the ID of *26* for the replication partner *array2*. The modification allows replication traffic Monday through Wednesday from 3 PM until 5 PM and changes the bandwidth.

```
Nimble OS $ partner --edit_throttle array2 --id 26
  --days Mon,Tue,Wed --at 15:00 --until 17:00
  --bandwidth 20 --description "afternoon replication limit"
```

This example deletes the throttle with the ID of 26.

```
Nimble OS $ partner --delete_throttle array1 --id 26
```

This example pauses all replication to the replication partner named *array2*.

```
Nimble OS $ partner --pause array2
```

This example resumes replication to the replication partner named *array2*.

```
Nimble OS $ partner --resume array2
```

This example tests connectivity to the replication partner named *array1*.

```
Nimble OS $ partner --test array1
Successfully contacted partner.
```

pe

The pe command options allow you to view protocol endpoints, which are administrative logical units (LU) in LU conglomerates to be used with VMware virtual volumes (VVOLs).

Synopsis

pe --help

pe --list

pe --info *name* [--verbose]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all protocol endpoints in the group.	
--info	*name*	Display detailed information about the specified protocol endpoint.	
	--verbose	Display additional information about the protocol endpoint, such as the composite LUN for each VVOL binding.	

perfpolicy

The perfpolicy command options manage performance policies.

You can use performance policies as a set of optimizations that includes block size, compression, deduplication, as well as caching. It ensure that the volume performance represents the best configuration for its intended use. By default, a volume uses the default performance policy, which is for 4096 byte blocks with full compression and caching enabled.

On creation, volumes inherit the dedupe default value of their associated Performance Policy unless explicitly overriden via the dedupe option. The dedupe attribute determines if new writes will be deduped. Upon cloning, the volume inherits the most recent dedupe settings of its parent.

Create and edit a performance policy assigning block size, Application Category, and default dedupe configuration. If an incorrect Application Category is specified, the error message will include a list of the valid Application Categories.

Block size should be set to match the application block size. If the application block size is not known, it should be set to the largest common denominator of the majority of I/O requests. If in doubt, set low (minimum 4KB), because setting it too high can reduce performance severely, while setting it too low reduces performance by a small amount.

Caching should generally be enabled. It helps most with non-sequential (random) access. Caching should be left on unless the volume is known to be either not read much or read sequentially (for example, a volume holding log files). If caching is turned off, the array will not cache any data for the volume but it continues to cache internal metadata.

Compression can, and should, be left on unless the volume is known to hold mostly incompressible data. In this case, turning off compression will improve performance by a small amount.

Synopsis

perfpolicy --help

perfpolicy --list

perfpolicy --info *policy_name*

perfpolicy --create *policy_name*
[--description *text*]
[--blocksize *bytes*]
[--compress *{yes|no}*]
[--cache *{yes|no}*]
[--cache_policy *{normal|aggressive}*]
[--space_policy *{offline|non_writable}*]
[--app_category *category*]
[--dedupe_enabled *{yes|no}*]

perfpolicy --edit *policy_name*

[--name *new_name*]
[--description *text*]
[--compress *{yes|no}*]
[--cache *{yes|no}*]
[--cache_policy *{normal|aggressive*]
[--space_policy *{offline|non_writable}*]
[--app_category *category*]
[--dedupe_enabled *{yes|no}*]

perfpolicy --delete *policy_name*
[--force]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all configured performance policies. Several predefined performance policies are provided, which are based on common application best practices.	Guest
--info	*policy_name*	Provide detailed information about the specified performance policy.	Guest

Option	Arguments and Suboptions	Description	Role
--create	*policy_name*	Create the specified performance policy. The block size of a performance policy cannot be changed once it has been created.	Operator
	[--description *text*]	Specify a description of the performance policy. If there are spaces in the description, enclose the text in quotation marks.	
	[--blocksize *bytes*]	Specify a block size to be used for volumes created with the specified performance policy. Supported block sizes are 4096 bytes (4 KB), 8192 bytes (8 KB), 16384 bytes (16 KB), and 32768 bytes (32 KB). If you do not specify this option, the default block size of 4096 bytes (4 KB) is used. You cannot change the after you create the performance policy.	
	[--compress *{yes|no}*]	Specify whether data in associated volumes should be compressed. If you do not specify this option, the default of yes is used, and data is compressed.	
	[--cache *{yes|no}*]	Specify whether data in associated volumes should be cached. If you do not specify this option, the default of yes is used, and data is cached.	
	[--cache_policy *{normal|aggressive}*]	Specify how data of an associated volume must be cached if data caching is enabled. The normal policy caches data but skips under certain conditions, such as sequential I/O. The aggressive policy accelerates caching of all data that is associated with this volume, regardless of sequentiality. You want to use the aggressive policy during the initial copy of a dataset for a new volume. However, it may cause the data that is associated with other volumes to be evicted from cache sooner. It may also cause greater wear on the flash-based cache over an extended period of time. Therefore, you may want to revert to normal caching after the data load has been completed. If you do not specify this option, the default of normal is used.	
	[--space_policy *{offline|non_writable}*]		

Option	Arguments and Suboptions	Description	Role
		Specify the state of the volume upon a space-constraint violation such as volume quota violation, snapshot quota violation, volumes above their volume reserve or snapshot reserve if the pool free space is exhausted. The "offline" space policy takes the volume offline. The "non_writable" space policy disallows write operations on the volume, but allows unmap and read operations. If you do not specify this option, the default of offline is used.	
	[--app_category *category*]	Specify the application category setting for volumes created with the specified performance policy. The application category is case sensitive. Valid application categories are: • Backup • DB2 • Exchange • File Server • Oracle • SharePoint • SQL Server • Virtual Desktop • Virtual Server • Other	
	[--dedupe_enabled *{yes | no}*]	Specify the dedupe enabled setting for volumes created with the specified performance policy. This default value can be overridden for each volume during creation or edit.	

Option	Arguments and Suboptions	Description	Role
--edit	*policy_name*	Modify the specified performance policy. If there are spaces in the *policy_name*, enclose the value in quotation marks. Specify at least one suboption.	Operator
	[--name *new_name*]	Modify the name of the performance policy	
	[--description *text*]	Modify the description of the performance policy.	
	[--compress *{yes|no}*]	Modify whether data in associated volumes should be compressed.	
	[--cache *{yes|no}*]	Modify whether data in associated volumes should be cached.	
	[--cache_policy *{normal|aggressive}*]	Modify the cache policy used when data in associated volumes is cached.	
	[--space_policy *{offline|non_writable}*]	Modify the space policy used.	
	[--app_category *category*]	Modify the application category setting for volumes created with the specified performance policy. The application category is case sensitive. Valid application categories are: • Backup • DB2 • Exchange • File Server • Oracle • SharePoint • SQL Server • Virtual Desktop • Virtual Server • Other	
	[--dedupe_enabled *{yes | no}*]	Modify the dedupe enabled setting for volumes created with the specified performance policy. This default value can be overridden for each volume during creation or edit.	
--delete	*policy_name*	Delete the specified performance policy. You cannot delete a performance policy that is associated with a volume. Remove the association before you delete the performance policy.	Power User
	[--force]	Forcibly delete the specified performance policy even if there are volumes associated with it.	

Examples

This example lists the configured performance policies, including predefined policies.

```
Nimble OS $ perfpolicy --list
--------------------------+----------+--------+-----+----------------------
Performance Policy Name   Block Size (bytes) Compress Cache Cache Policy
--------------------------+----------+--------+-----+----------------------
backup repository                 4096 Yes      Yes   normal
default                           4096 Yes      Yes   normal
Exchange 2003 data store          4096 Yes      Yes   normal
Exchange 2007 data store          8192 Yes      Yes   normal
Exchange 2010 data store         32768 Yes      Yes   normal
Exchange log                     16384 Yes      No    normal
Hyper-V CSV                       4096 Yes      Yes   normal
Oracle OLTP                       8192 Yes      Yes   normal
SharePoint                        8192 Yes      Yes   normal
SQL Server                        8192 Yes      Yes   normal
SQL Server 2012                   8192 Yes      Yes   normal
SQL Server Logs                   4096 Yes      No    normal
VMware ESX                        4096 Yes      Yes   normal
VMware ESX 5                      4096 Yes      Yes   normal
VMware VDI                        4096 Yes      Yes   normal
Windows File Server               4096 Yes      Yes   normal
```

This example shows details about the predefined performance policy named *default*.

```
Nimble OS $ perfpolicy --info default
Name: default
Description: Default performance policy
Block size: 4096 bytes
Compress: Yes
Cache: Yes
Cache Policy: normal
Space Policy: offline
Created: N/A
Last configuration change: N/A
```

This example creates a performance policy that specifies a block size of16384 bytes and includes compression, but does not use cache.

```
Nimble OS $ perfpolicy --create bigblock --description "big block, no cache"
--blocksize 16384 --cache no
```

This example modifies the performance policy named *bigblock* so that it uses cache.

```
Nimble OS $ perfpolicy --edit bigblock --cache yes
```

This example deletes performance policy after unassociating any associated volumes so that the --force option is not needed.

```
Nimble OS $ perfpolicy --delete bigblock
```

pool

Synopsis

The pool command options manage storage pools.

Pools can consist of either All-Flash Arrays or Adaptive Flash Arrays, but cannot have a mixture of array types.

pool --help

pool --list

pool --info *pool_name*

pool --create *pool_name*
--array *array*
[--description *description*]
[--dedupe_all_volumes {*yes*|*no*}]

pool --delete *pool_name*
[--force]

pool --edit *pool_name*
[--description *description*]
[--name *new_name*]
[--dedupe_all_volumes {*yes*|*no*}]
[--set_default_pool]

pool --assign *pool_name*
--array *array_name*
--allow_lower_limits

pool --unassign *pool_name*
--array *array*
[--force]

pool --merge *pool_name*
--target *target*
--allow_lower_limits

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all storage pools in the group.	Guest
--info	*pool_name*	Show detailed information about the specified storage pool. The default pool is assumed.	Guest

Option	Arguments and Suboptions	Description	Role
--create	*pool_name*	Create a storage pool.	Power User
	--array *array_name*	Specify which array to add to the storage pool. To add multiple arrays, repeat this option as many times as needed.	
	[--description *description*]	Specify a plain-text description of the storage pool. If there are spaces in the description, enclose the text in quotation marks.	
	[--dedupe_all_volumes {*yes*|*no*}]	Enables deduplication, by default, for all new volumes created. It does not affect the deduplication setting of existing volumes. Deduplication settings for existing volumes can be changed by editing the volume attributes.	
--delete	*pool_name*	Delete the specified storage pool. You cannot delete the default pool. Before you delete a pool, make sure that there are no volumes assigned to the pool. Use vol --list --pool *pool_name* to see if the pool you plan to delete has any volumes assigned to it. If any volumes are assigned, use vol --move to move those volumes to another storage pool. The default pool cannot be deleted.	Power User
	[--force]	Forcibly delete the specified storage pool, even if it contains volumes that were deleted but whose space is still being reclaimed.	
--edit	*pool_name*	Modify the storage pool settings. Specify at least one suboption.	Power User
	[--description *text*]	Modify the plain-text description of the storage pool.	
	[--name *new_name*]	Modify the name of the storage pool.	
	[--dedupe_all_volumes {*yes*|*no*}]	Enable or disable deduplication for all new volumes created. This does not affect the deduplication setting of existing volumes. Deduplication settings for existing volumes can be changed by editing the volume attributes.	
	[--set_default_pool]	Make this pool the default. Once this is set, future operations that do not explicitly specify a pool will use this pool	

Option	Arguments and Suboptions	Description	Role
--assign	*pool_name*	Assign arrays to the storage pool.	Power User
	--array *array*	Specify an array to add to the pool. To assign multiple arrays to the pool, repeat this option as many times as needed.	
	--allow_lower_limits	Allow pool to have lower limits after the operation. When an array is assigned to an existing pool, and the pool limits are higher than the array limits, the resulting pool will have the lower limits of the array, which could also lower the group limits.	
--unassign	*pool_name*	Unassign, or remove, arrays from the specified storage pool.	Power User
	--array *array*	Specify an array to remove from the pool. To unassign multiple arrays from the pool, repeat this option as many times as needed.	
	[--force]	Forcibly unassign the array or arrays from the pool even if the arrays are unreachable. For this operation to succeed, there should be no volumes associated with the array pool.	
--merge	*pool_name*	Merge the specified storage pool into the specified target pool. All volumes on the specified pool are moved to the target pool and the specified pool is then deleted. Likewise, all arrays in the pool are assigned to the target pool. The name of the newly merged pool is automatically truncated to a maximum of 255 characters.	Power User
	--target *target*	Specify the target, or destination, pool into which the storage pool is merged.	
	--allow_lower_limits	When merging two pools with different limits, the resulting pool will have the limits of the smallest pool prior to the merge.	

Examples

This example lists all configured storage pools. In this example, only the default storage pool appears.

```
Nimble OS $ pool --list
--------------+-----------+----------+----------------------------------------
Pool          Capacity(MB) Usage(MB)  Arrays
--------------+-----------+----------+----------------------------------------
default           15738305    2719551 array7
```

This example shows detailed information about the default group.

```
Nimble OS $ pool --info default
Pool Name: default
Pool Description: Default pool
Created: N/A
Last configuration change: N/A
Arrays in the pool:
Array mktg, Capacity (MB): 185502916, Usage (MB): 11468728
Array sales, Capacity (MB): 65394722, Usage (MB): 8399605
Arrays being unassigned from the pool:
Pool capacity (MB): 250897638
Pool cache capacity (MB): 305152
Pinnable cache capacity (MB): 662
Pinned cache capacity (MB): 0
Pinned volumes :
Pool usage (MB): 19868333
Pool unused reserve (MB): 20482561
Pool free space (MB): 210546744
Volumes in the pool: vol10,vol11,SQL-2012-Server,SQL-2012-Server-Data,
SQL-2012-Server-Log
Snapshot collection count: 1200
Snapshot count: 7092
```

This example creates a storage pool named *marketing* that includes the array named array7.

```
Nimble OS $ pool --create marketing --array array7
```

This example changes the name of the *marketing* storage pool to *marketingteam*.

```
Nimble OS $ pool --edit marketing --name martketingteam
```

This example deletes a storage pool named *marketingteam*.

```
Nimble OS $ pool --delete marketingteam
```

This example assigns an array named *MKT2* to the storage pool named *marketingteam*.

```
Nimble OS $ pool --assign marketingteam --array MKT2
```

This example unassigns an array named *MKT2* from the storage pool named *marketingteam*.

```
Nimble OS $ pool --unassign marketingteam --array MKT2
```

This example merges the storage pool named *marketingteam* into the storage pool named *corporatepool*.

```
Nimble OS $ pool --merge marketingteam --target corporatepool
```

prottmpl

The prottmpl command options manage protection templates, on which volume collection schedules and retention are based. You cannot edit or delete the protection templates provided by Nimble Storage.

Protection templates provide a way to create a set of data protection policies that meet specific goals, and apply them as volumes and collections are created. Protection templates consist of sets of snapshot schedules, replication schedules, and retention limits that can be used to prefill the protection information when creating volume collections. A volume collection, once created, us nit affected by edits to the protection template that was used to create it, all the volumes assigned to a volume collection use the same settings.

You cannot edit or delete the predefined protection templates provided with the NimbleOS software, but you can create custom protection templates as needed.

Synopsis

prottmpl --help

prottmpl --list

prottmpl --info *name*

prottmpl --create *name*
[--description *text*]
--app_sync {*none|vss|vmware|generic*}
[--app_server *server*]
[--app_id {*exchange|sql2005|sql2008|sql2012|exchange_dag|sql2014|sql2016|hyperv*}]
[--app_cluster_name *cluster_name*]
[--app_service_name *service_name*]
[--vcenter_hostname *server*]
[--vcenter_username *server*]
[--vcenter_password *server*]
[--agent_hostname *server*]
[--agent_username *user_name*]
[--agent_password *password*]
--schedule *name*
--repeat *period*
--repeat_unit {*minutes|hours|days|weeks*}
[--at *time*]
[--until *time*]
[--days *days*]
[--retain *number*]
[--replicate_to *partner*]
[--replicate_every *number*]
[--num_retain_replica *number*]
[--alert_threshold *hh:mm*]
[--snap_verify {*yes|no*}]
[--skip_db_consistency_check {*yes|no*}]
[--disable_appsync {*yes|no*}]
[--external_trigger {*yes|no*}]

prottmpl --delete *name*

prottmpl --edit *name*
[--newname *name*]
[--description *text*]
[--app_sync {*none|vss|vmware|generic*}]
[--app_server *server*]
[--app_id {*exchange|sql2005|sql2008|sql2012|exchange_dag|sql2014|sql2016|hyperv*}]
[--app_cluster_name *cluster_name*]
[--app_service_name *service_name*]
[--vcenter_hostname *server*]
[--vcenter_username *server*]
[--vcenter_password *server*]
[--agent_hostname *server*]
[--agent_username *user_name*]
[--agent_password *password*]

prottmpl --addsched *name*
--schedule *name*
--repeat *period*
--repeat_unit {*minutes|hours|days|weeks*}
[--at *time*]
[--until *time*]
[--days *days*]
[--retain *number*]
[--replicate_to *partner*]
[--replicate_every *number*]
[--num_retain_replica *number*]
[--alert_threshold *hh:mm*]
[--snap_verify {*yes|no*}]
[--skip_db_consistency_check {*yes|no*}]
[--disable_appsync {*yes|no*}]
[--external_trigger {*yes|no*}]

prottmpl --editsched *name*
--schedule *name*
[--name *new_name*]
[--repeat *period*]
[--repeat_unit {*minutes|hours|days|weeks*}]
[--at *time*]
[--until *time*]
[--days *days*]
[--retain *number*]
[--replicate_to *partner*]
[--replicate_every *number*]
[--num_retain_replica *number*]
[--alert_threshold *hh:mm*]
[--snap_verify {*yes|no*}]
[--skip_db_consistency_check {*yes|no*}]
[--disable_appsync {*yes|no*}]
[--external_trigger {*yes|no*}]

prottmpl --deletesched *template_name*
--schedule *schedule_name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List protection templates.	Guest
--info	*name*	Provide detailed information about the specified protection template and its schedules.	Guest

Option	Arguments and Suboptions	Description	Role
--create	*name*	Create a protection template with the given schedule.	Operator
	[--description *text*]	Specify a plain-language description of the protection template. This is useful to specify what the protection template was designed to optimize. If there are spaces in the name, enclose the text in quotation marks,	
	--app_sync {*none*|*vss*|*vmware*|*generic*}	Define the application synchronization to apply. If you specify *vss* synchronization, provide the --app_* credentials. If you specify *vmware* synchronization, provide the --vcenter_* credentials. If you have created a custom agent, specify the *generic* app type when creating a protection template.	
	[--app_server *server*]	Specify the application server hostname when *vss* synchronization is enabled.	
	[--app_id {*exchange*|*sql2005*|*sql2008*|*sql2012*|*exchange_dag*|*sql2014*|*sql2106*|*hyperv*}]	Define the application running on the server when *vss* synchronization is enabled, including MS Hyper-V.	
	[--app_cluster_name *cluster_name*]	Specify the cluster name if the application that uses *vss* synchronization is running within a Windows clustering environment.	
	[--app_service_name *service_name*]	Specify the instance name of the service if the application that uses *vss* synchronization is running within a Windows clustering environment.	
	[--vcenter_hostname *server*]	Specify the host hostname of the vCenter server with which the Nimble array will communicate when *vmware* synchronization is enabled. A customer port number can be specified using ":".	
	[--vcenter_username *server*]	Specify the Windows user name to use to create a Nimble account on the vCenter server when *vmware* synchronization is enabled.	
	[--vcenter_password *server*]	Specify the password associated with the user name when *vmware* synchronization is enabled.	

Option	Arguments and Suboptions	Description	Role
	[--agent_hostname *server*]	Backup hostname. A custom port number can be specified using ":".	
	[--agent_username *user_name*]	Backup agent user name.	
	[--agent_password *password*]	Backup agent user's password.	
	[--schedule *name*]	Name of the snapshot schedule to associate with this protection template. Use --addsched to add more schedules to the protection template after creating it.	
	--repeat *period*	Specify the frequency of snapshots every --repeat_unit. If you do not specify this option, the default value of *1* is used.	
	--repeat_unit {*minutes*|*hours*|*days*|*weeks*}	Specify the unit of time for --repeat to define the frequency of snapshots. If you do not specify this option, the default value of to *days* is used.	
	[--at *time*]	Specify the time of day to start taking snapshots. If you do not specify this option, the default value of *12:00* (midnight) is used.	
	[--until *time*]	Specify the time of day to stop taking snapshots. If you do not specify this option, the default value of *11:59* (PM) is used if a stop time is appropriate for the repeat unit. For example, if you take snapshots once daily, you do not need this option.	
	[--days *days*]	Specify on which days to take snapshots unless the repeat unit is *weeks*. You can enter a comma-separated list of days, such as Mon,Wed,Fri. If you do not specify this option, the default value of *all* is used.	
	[--retain *number*]	Specify the maximum number of snapshots to keep. After this number is reached, older snapshots can be deleted to make room for new ones if the space is needed. If replication is enabled on this schedule, the array always retains the latest replicated snapshot, which may exceed the specified retention value. This is necessary to ensure efficient replication performance.	

Option	Arguments and Suboptions	Description	Role
	[--replicate_to *partner*]	Specify the replication partner for replicated snapshots. Use partner --create on both arrays to configure replication partners.	
	[--replicate_every *number*]	Specify that a certain number of snapshots assigned to the schedule should be replicated. For example, setting this to 5 replicates every fifth snapshot. If snapshots are replicated and you do not specify this option, all snapshots are replicated.	
	[--num_retain_replica *number*]	Specify the number of snapshots to retain on the replication partner.	
	[--alert_threshold *hh:mm*]	If replicating a snapshot takes more than this amount of time to complete, an alert will be generated. If you do not specify this option, the default of 24 hours is used. Enter 00:00 to disable this alert.	
	[--snap_verify {*yes*|*no*}]	Specify whether to run a verification tool on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization. The tool used to verify snapshots depends on the type of application. For example, if the application synchronization is *vss* and the application ID is *exchange*, the eseutil tool is run on the snapshots. If verification fails, the logs are not truncated.	
	[--skip_db_consistency_check {*yes*|*no*}]	Specify whether to skip consistency checks for database files on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization of *vss*, the application ID is *exchange_dag*. (MS Exchange 2010 or later using Database Available Group), --snap_verify is set to *yes*, and --disable_appsync is set to *no*. Skipping consistency checks is only recommended if each database in a DAG has multiple copies.	
	[--disable_appsync {*yes*|*no*}]		

Option	Arguments and Suboptions	Description	Role
		Specify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	
	[--external_trigger {*yes*|*no*}]	Specify whether this schedule is managed by externally driven manual snapshots. Normal protection schedules have internal timers which drive snapshot creation. An externally drive schedule has no internal timers, and all snapshot activity is driven by an external trigger.	

Option	Arguments and Suboptions	Description	Role
--edit	*name*	Modify attributes of the specified protection template. New volume collections based on the template are affected by the modifications, but existing volume collections are not. Specify at least one suboption.	Operator
	[--name *new_name*]	Modify the name of the protection template.	
	[--description *text*]	Modify the plain-language description of the protection template.	
	--app_sync {*none*|*vss*|*vmware*|*generic*}	Modify the application synchronization to apply. If you have created a custom agent, specify the *generic* app type when editing a protection template.	
	[--app_server *server*]	Modify the application server for *vss* synchronization.	
	[--app_id {*exchange*|*sql2005*|*sql2008*|*sql2012*|*exchange_dag*|*sql2014*|*sql2016*|*hyperv*}]	Modify the application running on the server for *vss* synchronization, including MS Hyper-V.	
	[--app_cluster_name *cluster_name*]	Modify the cluster name if the application that uses *vss* synchronization is running within a Windows clustering environment.	
	[--app_service_name *service_name*]	Modify the instance name of the service if the application that uses *vss* synchronization is running within a Windows clustering environment.	
	[--vcenter_hostname *server*]	Modify the host name of the vCenter server with which the Nimble array communicates for *vmware* synchronization.	
	[--vcenter_username *server*]	Modify the Windows user name for *vmware* synchronization.	
	[--vcenter_password *server*]	Modify the password associated with the user name for *vmware* synchronization.	
	[--agent_hostname *server*]		
	[--agent_username *user_name*]		
	[--agent_password *password*]		
--delete	*name*	Delete the specified protection template.	Operator

Option	Arguments and Suboptions	Description	Role
--addsched	*name*	Add a schedule to the specified protection template.	Operator
	--schedule *name*	Specify the name of the snapshot schedule to associate with this protection template.	
	--repeat *period*	Specify the frequency of snapshots. If you do not specify this option, the default value of *1* is used.	
	--repeat_unit {*minutes*|*hours*|*days*|*weeks*}	Specify the unit of time for to define the frequency of snapshots. If you do not specify this option, the default value of *days* is used.	
	[--at *time*]	Specify the time of day to start taking snapshots. If you do not specify this option, the default value of *12:00* (midnight) is used.	
	[--until *time*]	Specify the time of day to stop taking snapshots. If you do not specify this option, the default value of *11:59* (PM) is used if a stop time is appropriate for the repeat unit. For example, if you take snapshots once daily, you do not need this option.	
	[--days *days*]	Specify the days on which to take snapshots unless the repeat unit is *weeks*. You can enter a comma-separated list of days, such as Mon,Wed,Fri. If you do not specify this option, the default value of *all* days of the week is used.	
	[--retain *number*]	Specify the maximum number of snapshots to keep. After this number is reached, older snapshots can be deleted to make room for new ones if the space is needed. If replication is enabled on this schedule, the array always retains the latest replicated snapshot, which may exceed the specified retention value. This is necessary to ensure efficient replication performance.	
	[--replicate_to *partner*]	Specify the replication partner for replicated snapshots. Use partner --create on both arrays to configure replication partners.	
	[--replicate_every *number*]		

Option	Arguments and Suboptions	Description	Role
		Specify that a certain number of snapshots assigned to the schedule should be replicated. For example, setting this to 5 replicates every fifth snapshot. If snapshots are replicated and you do not specify this option, all snapshots are replicated.	
	[--num_retain_replica *number*]	Specify the number of snapshots to retain on the replication partner.	
	[--alert_threshold *hh:mm*]	If replicating a snapshot takes more than this amount of time to complete, an alert will be generated. If you do not specify this option, the default of 24 hours is used. Enter 00:00 to disable this alert.	
	[--snap_verify {*yes*|*no*}]	Specify whether to run a verification tool on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization. The tool used to verify snapshots depends on the type of application. For example, if the application synchronization is *vss* and the application ID is *exchange*, the eseutil tool is run on the snapshots. If verification fails, the logs are not truncated.	
	[--skip_db_consistency_check {*yes*|*no*}]	Specify whether to skip consistency checks for database files on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization of *vss*, the application ID is *exchange_dag*. (MS Exchange 2010 or later using Database Available Group), --snap_verify is set to *yes*, and --disable_appsync is set to *no*. Skipping consistency checks is only recommended if each database in a DAG has multiple copies.	
	[--disable_appsync {*yes*|*no*}]	Specify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	
	[--external_trigger {*yes*|*no*}]]		

Option	Arguments and Suboptions	Description	Role
		Specify whether this schedule is managed by externally driven manual snapshots. Normal protection schedules have internal timers which drive snapshot creation. An externally drive schedule has no internal timers, and all snapshot activity is driven by an external trigger.	

Option	Arguments and Suboptions	Description	Role
--editsched	*name*	Modify a schedule for the specified protection template. Specify at least one suboption in addition to --schedule *name*.	Operator
	--schedule *name*	Specify the name of the snapshot schedule to modify.	
	[--newname *name*]	Modify the name of the snapshot schedule.	
	[--repeat *period*]	Modify the frequency of snapshots for each --repeat_unit.	
	[--repeat_unit {*minutes|hours|days|weeks*}]	Modify the unit of time for --repeat to define the frequency of snapshots.	
	[--at *time*]	Modify the time of day to start taking snapshots.	
	[--until *time*]	Modify the time of day to stop taking snapshots.	
	[--days *days*]	Modify the days on which to take snapshots unless the repeat unit is *weeks*.	
	[--retain *number*]	Modify the maximum number of snapshots to keep.	
	[--replicate_to *partner*]	Modify the replication partner for replicated snapshots.	
	[--replicate_every *number*]	Modify that a certain number of snapshots assigned to the schedule should be replicated.	
	[--num_retain_replica *number*]	Modify the number of snapshots to retain on the replication partner.	
	[--alert_threshold *hh:mm*]	Modify whether to generate an alert if replicating a snapshot takes more than this amount of time to complete.	
	[--snap_verify {*yes|no*}]	Modify whether to run a verification tool on snapshots created by this schedule.	
	[--skip_db_consistency_check {*yes|no*}]	Modify whether to skip consistency checks for database files on snapshots created by this schedule.	
	[--disable_appsync {*yes|no*}]	Modify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	

Option	Arguments and Suboptions	Description	Role
	[--external_trigger {*yes*|*no*}]	Modify whether this schedule is managed by externally driven manual snapshots. Normal protection schedules have internal timers which drive snapshot creation. An externally drive schedule has no internal timers, and all snapshot activity is driven by an external trigger.	
--deletesched	*name*	Delete the specified schedule from the specified protection template.	Operator
	--schedule *name*	Specify the snapshot schedule to delete.	

Examples

This example lists available protection templates. In this example, only the predefined templates exist.

```
Nimble OS $ prottmpl --list
-----------------------------------+-----------------------------------------
Protection Template                Application
Name                               Synchronization
-----------------------------------+-----------------------------------------
Retain-30Daily                     none
Retain-48Hourly-30Daily-52Weekly   none
Retain-90Daily                     none
```

This example shows detailed information about the *Retain-30Daily* protection template.

Note: This protection template is one of the predefined templates, so you cannot modify or delete it.

```
Nimble OS $ prottmpl --info Retain-30Daily
Name: Retain-30Daily
Description: Provides daily snapshots retained for 30 days
Application synchronization: none
Application server: N/A
Application ID: N/A
Cluster name: N/A
Service name: N/A
VMware vCenter hostname: N/A
VMware vCenter username: N/A
VMware vCenter password: N/A
Created: N/A
Last configuration change: N/A
Schedule Name: daily
        Description:
        Type: Nimble Schedule
        Repeats: 1 day(s)
        At: 0:00:00 a.m.
        Until: N/A
        Days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
        Number of snapshots to retain: 30
        Snapshot verification: N/A
        Skip database consistency check: N/A
        Disable appsync: No
```

```
        Number of snapshots to retain on replica: 0
        Replicate every: N/A
        Replicate to:
        Alert threshold: 24:00
        Created: N/A
        Last configuration change: N/A
        Last timing configuration change: N/A
```

This example creates a new protection template named *highrepl*.

```
Nimble OS $ prottmpl --create highrepl
--description "use when high replication is needed"
--app_sync none --schedule replicated --repeat 1 --repeat_unit day
--days Monday,Tuesday,Thursday --retain 8 --replicate_to array7
--replicate_every 1 --num_retain_replica 40 --snap_verify no
```

This example modifies the protection template named *highrepl* to replicate once a week, and to retain 31 snapshots and 52 replicas.

```
Nimble OS $ prottmpl --edit highrepl --schedule replicated
--repeat 1 --repeat_unit weeks --retain 31 --replicate_to array7
--replicate_every 1 --num_retain_replica 52 --snap_verify no
```

This example deletes a protection template named *highrepl*.

```
Nimble OS $ prottmpl --delete highrepl
```

This example adds an *hourly* schedule to the protection template named *highrepl*.

```
Nimble OS $ prottmpl --addsched highrepl
--schedule hourly --repeat 1 --repeat_unit hours
--retain 25  --snap_verify no
```

This example modifies the *hourly* schedule that is associated with the protection template named *highrepl*.

```
Nimble OS $ prottmpl --editsched highrep  --schedule hourly
--repeat 2 --repeat_unit hours
```

This example removes a schedule named *hourly* from the protection template named *highrepl*.

```
Nimble OS $ prottmpl --deletesched highrepl  --schedule hourly
```

This example shows creating a protection template with app sync type generic.

```
Nimble OS $ prottmpl --create gen-sync-test-prottmpl --app_sync generic
--agent_hostname 10.18.237.121 --agent_username Administrator --agent_password
 Nim123# --schedule sched1 --retain 30
```

This example shows editing a protection template with app sync type generic.

```
prottmpl --edit gen-sync-test-prottmpl --agent_hostname 10.xx.xx.xxx
```

reboot

This command reboots an entire group of arrays, an array, or a controller. If you want to stop I/O traffic without rebooting, you must use the halt command. The reboot command causes a reboot without a power cycle.

Synopsis

reboot --help

reboot --array *array_name*
[--controller *{A|B|*]
[--non_interactive]

reboot --group
[--non_interactive]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--array	*array_name*	Reboot the specified array.	Power User
	[--controller (*A|B*)]	Reboot the specified controller on an array. Valid controller names are *A* and *B*, which are case sensitive. If you do not specify this option, both controllers on the specified array are rebooted.	
	--non_interactive	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly taking the array or controller offline for a short time and disrupting client connections.	
--group	N/A	Reboot the entire group when you run this command on the group leader. If any array fails to respond to the reboot request, then the group leader array remains operational.	Power User
	--non_interactive	Override the default behavior of having to type "yes" to complete the command. Use this suboption with care to avoid unexpectedly taking the group offline for a short time and disrupting client connections.	

Examples

This example cancels rebooting the array where you are logged in.

```
Nimble OS $ reboot
WARNING: This operation will reboot the array.
```

```
Type yes to continue, no to cancel: no
INFO: Canceling reboot
```

This example reboots the group without prompting for confirmation. You must run this command on the group leader array.

```
Nimble OS $ reboot --group --non_interactive
```

route

The route command options manage network routes. Specify a network route either in "network/netmask" (dotted quad) format (for example, 10.12.0.100/255.255.0.0) or "network/mask_bits" (slash) format (for example, 10.12.0.100/16). When the system must connect to a remote array or host, it checks the routing table to determine if a known path exists. If the remote system is in a subnet to which the system has access, it tries to connect along that interface.

Synopsis

route --help

route --list
[--netconfig *{active|backup|draft}*]

route --info *network*
[--netconfig *{active|backup|draft}*]
[--gateway *gateway*

route --add *network*
--netconfig *{active|backup|draft}*
--gateway *gateway*

route --edit *network*
[--netconfig *{active|backup|draft}*]
[--network *network*]
[--gateway *gateway*]

route --delete *network*
[--netconfig *{active|backup|draft}*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List all configured routes.	Guest
	[--netconfig *{active|backup|draft}*]	List all configured routes on the specified network configuration.	

Option	Arguments and Suboptions	Description	Role
--info	*network*	Provide detailed information about the specified network route. The route is Destination/Netmask from the --list output.	Guest
	[--netconfig *{active|backup|draft}*]	Provide detailed information about the specified network route on the specified network configuration.	
	[--gateway*gateway*]	Provide detailed information about the network route for the specified gateway.	
--add	*network*	Add the specified network route.	Power User
	--netconfig *{active|backup|draft}*	Add the specified network route on the specified network configuration.	
	--gateway *gateway*	Add the specified network route for the specified gateway.	
--edit	*network*	Modify the specified network route. Specify at least one suboption.	Power User
	[--netconfig *{active|backup|draft}*]	Modify the specified network route on the specified network configuration.	
	[--network *network*]	Modify the specified network route for a new destination network.	
	[--gateway *gateway*]	Modify the specified network route for the specified gateway.	
--delete	*network*	Delete the specified network route.	Power User
	[--netconfig *{active|backup|draft}*]	Specify the network configuration from which to delete the specified network route.	

Examples

This example lists the routes in the *active* network configuration.

```
Nimble OS $ route --list --netconfig active
---------------+---------------+---------------
Destination     Netmask         Gateway
---------------+---------------+---------------
0.0.0.0         0.0.0.0         10.12.12.1
10.1.0.0        255.255.255.0   10.11.0.2
```

This example provides detailed information about one route. Notice how the *network* value is Destination/Netmask from the --list output.

```
Nimble OS $ route --info 10.1.0.0/255.255.255.0
Destination: 10.1.0.0
Netmask: 255.255.255.0
Gateway: 10.11.0.2
```

This example adds a network route to the *draft* network configuration.

```
 route --add 10.12.134.209/255.255.0.0 --gateway 0.12.0.0
--netconfig draft
```

This example edits a network route.

```
route --edit 0.12.134.209/255.255.0.0 --network 10.12.136.209/255.255.0.0
--gateway 10.12.0.0
```

This example deletes the route at 0.0.0.0/0.0.0.0.

```
Nimble OS $ route --delete 0.0.0.0/0.0.0.0
```

setup

The setup command and its options configure an array.

You can set up an array with the interactive setup wizard in the command line interface. To learn more about using the interactive setup wizard, see the *Hardware Guide*. Alternatively, you can specify the options in the CLI as described here. If you use the non-interactive command line, then you must specify all mandatory settings as command line options.

Synopsis

setup --help

setup
[--group_name *group_name*]
--name *array_name*
[--domainname *domain_name*]
[--dnsserver *server*]
[--ntpserver *server*]
[--timezone *zone*]
--mgmt_ipaddr *ipaddr*]
--backup_mgmt_ipaddr *ipaddr*]
--subnet_label *label*
--subnet_addr *network/netmask*
[--subnet_type {*{mgmt|data|mgmt,data*}]
[--subnet_mtu *mtu*]
--default_gateway *ipaddr*
[--data_ipaddr *ipaddr*]
[--discovery_ipaddr *ipaddr*]
[--support_ipaddr *ipaddr*]
[--iscsi_automatic_connection_method {*yes|no*}]
[--iscsi_connection_rebalancing {*yes|no*}]
--accept_license
[--view_license]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
	[--group_name *group_name*]	Specify the name of a group. A single-array group is by default named the *default* group.	Administrator
	--name *array_name*	Specify the name of the array.	
	[--domainname *domain_name*]	Specify a DNS domain name.	
	[--dnsserver *server*]	Specify the IP address of the DNS server to use. Repeat this option to specify IP addresses for multiple DNS servers.	
	[--ntpserver *server*]	Specify a host name or an IP address for the NTP server to use.	
	[--timezone *zone*]	Specify the time zone the array is in.	
	--mgmt_ipaddr *ipaddr*	Specify a management IP address.	
	--backup_mgmt_ipaddr *ipaddr*	Specify a backup management IP address.	
	--subnet_label *label*	Specify a subnet on a NIC. This option can be repeated to specify more than one subnet label. The first subnet label corresponds to the first NIC, vizz, or eth1, and the second subnet label is assigned to the second NIC, viz, or eth2, and so on.A blank argument can be used to skip one NIC.	
	--subnet_addr*network/netmask*	Specify the subnet IP address. This option can be repeated to specify more than one subnet. The first subnet is assigned to the first NIC, viz., eth1, and the second subnet is assigned to the second NIC, viz1, or eth2, and so on. A blank argument can be used to skip one NIC.	
	[--subnet_type {*mgmt|data|mgmt,data*}]	Specify the subnet type. Valid subnet types are management and data. A subnet can be both management and data. Multiple subnet types can be specified by using a comma separated list. This option can be repeated to specify more that one subnet type. The first subnet type is assigned to the first NIC via eth1, and the second subnet is assigned to the second NIC via eth2, and so on. A blank argument can be used to skip one NIC. In a single-array Fibre Channel group, a data subnet is optional. Use the management subnet for replication.	
	[--subnet_mtu *mtu*]	Specify the subnet maximum transmission unit (MTU). If you do not specify this option, the default of *1500* is used.	

Option	Arguments and Suboptions	Description	Role
	--default_gateway *ipaddr*	Specify the default gateway IP address.	
	[--data_ipaddr *ipaddr*]	Specify the data IP address. This option can be repeated to specify more than one NIC IP address. The first IP address is assigned to the first NIC via eth1, and the second IP address is assigned to the second NIC via eth2, and so on. Blank argument can be used for skipping one NIC. Configure at least on data IP on iSCSI arrays. On Fibre Channel arrays, you can specify an IP address, or leave the value empty (--data_ipaddr '') to bypass this option. If IP addresses are specified on Fibre Channel arrays, the data IPs are used to handle replication traffic and entrap-group communication.	
	[--discovery_ipaddr *ipaddr*]	Specify the discovery IP address. If you do not specify this option, the management IP address is used.	
	[--support_ipaddr *ipaddr*]	Specify the support IP address. Include this argument twice. Then the first IP address entered is assigned to Controller A and the second to Controller B. You must assign static IP addresses to allow array access if the management IP address is not accessible.	
	[--iscsi_automatic_connection_method {*yes*|*no*}]	Redirect connections to the best data IP address based on connection counts. If you do not specify this option on an iSCSI array, the default of *yes* is used. This option has no effect in a Fibre Channel array.	
	[--iscsi_connection_rebalancing {*yes*|*no*}]	Rebalance the iSCSI connection by periodically breaking existing connections that are out of balance, which allows the host to reconnect to a more appropriate data IP. If you do not specify this option on an iSCSI array, the default of *yes* is used. This option has no effect in a Fibre Channel array.	
	--accept_license	Accept the terms and conditions of use, the end user license agreement (EULA), and third-party software notices.	
	[--view_license]	List the terms and conditions of use, the EULA, and third-party software notices.	

Examples

This example non-interactively sets up an iSCSI array.

```
Nimble OS $ setup --group_name default \
--name array7 \
--domainname admin.nimblestorage.com \
--dnsserver 8.8.8.8 \
--ntpserver time.nimblestorage.com \
--timezone America/Los_Angeles
--subnet_addr 10.12.147.253/255.255.0.0 --subnet_type management,data \
--subnet_addr 10.12.155.253/16 --subnet_type management,data \
--subnet_addr 10.12.0.0/255.255.0.0 --subnet_type management,data \
--subnet_addr 10.12.0.0/16 --subnet_type management,data \
--mgmt_ipaddr 10.12.147.253 \
--discovery_ipaddr 10.12.145.253 \
--support_ipaddr 10.12.149.253 \
--support_ipaddr 10.12.151.253 \
--data_ipaddr 10.12.153.253 \
--data_ipaddr 10.12.155.253 \
--data_ipaddr 10.12.157.253 \
--data_ipaddr 10.12.159.253 \
--default_gateway 10.12.255.254 \
--accept_license
```

shelf

The shelf command options manage expansion shelves on an array. An array can be a standalone head shelf, or a head shelf with one to seven shelves connected to it, depending on the array and shelf models.

Synopsis

shelf --help

shelf --list
--array *name*
[--ctrlr {*A*|*B*}]

shelf --info *shelf_serial*
--array *name*
[--verbose]

shelf --show_topology
--array *name*
[--ctrlr {*A*|*B*}]

shelf --activate *shelf_serial*
--array *name*

[--driveset *index*]
[--accept_foreign]
[--accept_dedupe_impact]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	--array *name*	List all shelves connected to an array or to a specified controller, A or B. Information for each shelf includes its serial number, model number, state, location and sensor status.	Guest
	[--ctrlr {*A*|*B*}]	List expansion shelves on the specified controller on the specified array.	
--info	*shelf_serial* --array *name*	Show information about a specific expansion shelf by serial number. Get the *shelf_serial* from --list output. Use the --array option to display information about all shelves on the array.	Guest
	[--verbose]	Show detailed information about a specific expansion shelf.	
--show_topology	--array *name*	List the shelves on an array and show how they are connected to the array or to a specific controller, A or B.	Guest
	[--ctrlr {*A*|*B*}]	Show the topology for all shelves connected to the specified controller.	
--activate	*shelf_serial* --array *name*	Activate the specified expansion shelf on the specified array. Activating an expansion shelf prepares and formats it to store user data. Before activating a shelf, be sure that all shelf cables are properly connected and the shelf is powered on.	Power User
	[--driveset *index*]	Activate a specific drive set, 0 (first drive set) or 1 (second drive set).	
	[--accept_foreign]	Accept the removal of data on the foreign shelf and activate it.	
	[--accept_dedupe_impact]	Accept the reduction or elimination of deduplication capabilities on the array as the result of activating a shelf that does not meet the necessary deduplication requirements.	

Examples

This example lists the expansion shelves on an array named *array7*.

```
Nimble OS $ shelf --list --array array7
-----------+---------+-------+---------+-------+-------+-------+---------
Serial #    Model     State   Location Power    Fans    Temp    Is Flash
                                       Supply   Status  Status  Shelf
                                       Status
-----------+---------+-------+---------+-------+-------+-------+---------
```

```
AC-109019   CS700     online   A.0      OK       OK      OK      No
AC-103333   ES1-H65   online   A.P1.1   OK       OK      OK      No
AC-103322   ES1-H65   online   A.P2.1   OK       OK      OK      No
AC-109019   CS700     online   B.0      OK       OK      OK      No
AC-103333   ES1-H65   online   B.P1.1   OK       OK      OK      No
AC-103322   ES1-H65   online   B.P2.1   OK       OK      OK      No
```

This example shows detailed information about one expansion shelf on an array named *array7*. The output is truncated to save space.

```
Nimble OS $ shelf --info AC-109019 --array array7
Serial #: AC-109019
Model: CS700
25922 2014-09-05,18:56:52.882370-07 ERROR: plat:_plat_get_hw_nvram_card_type:
 Failed to open /proc/bus/pci/devices (errno 2)
Extended Model: CS700-2T4F-48T-3200FS
Is Flash Shelf: No
Location: A.0 B.0
State: online
Usable capacity (MB): 32001478
Raw capacity (MB): 45785374
Usable cache capacity (MB): 3052288
Raw cache capacity (MB): 3052391
Power supply: OK
        power-supply1 at left rear: OK
        power-supply2 at right rear: OK
Cooling fans: OK
        fan1-A at lower front of controller A: 11772 rpm
        fan2-A at lower left rear of controller A: 11908 rpm
        fan3-A at lower right rear of controller A: 12223 rpm
        fan4-A at upper right front of controller A: 11670 rpm
        fan5-A at upper left front of controller A: 12053 rpm
        fan6-A at upper left rear of controller A: 12061 rpm
        fan1-B at lower front of controller B: 9486 rpm
        fan2-B at lower left rear of controller B: 10242 rpm
        fan3-B at lower right rear of controller B: 10072 rpm
        fan4-B at upper right front of controller B: 9758 rpm
        fan5-B at upper left front of controller B: 9851 rpm
        fan6-B at upper left rear of controller B: 9928 rpm
Temperature sensors: OK
        motherboard-A at motherboard: 38 degree C
        bp-temp1-A at left-side backplane: 24 degree C
        motherboard-B at motherboard: 38 degree C
        bp-temp2-B at right-side backplane: 20 degree C
Location on A: 0
Location on B: 0
Hardware state of A: ready
Hardware state of B: ready
Shelf sw type of A: Head Shelf
Shelf sw type of B: Head Shelf
Last known disk set:
  Number of disks: 16
  ...
  ...
```

This example shows the topology of expansion shelves on controller A of an array named *array7*.

```
Nimble OS $ shelf --show_topology --array array7 --ctrlr A
[AC-109019].A.SAS PORT1 --> [AC-103333].A.IN
[AC-109019].A.SAS PORT2 --> [AC-103322].A.IN
```

This example adds and activates an expansion shelf on an array named *array7*.

```
Nimble OS $ shelf --activate AA-102081 --array array7
```

snap

The snap command options manage snapshots.

Snapshots are point-in-time copies of volumes that can be brought online and used immediately, cloned, and replicated. Like volumes, snapshots can be accessed by initiators, are subject to the same controls, can be modified, and have the same restrictions.

The initial snapshot shares its original data with the volume from which it was created. Each successive snapshot consumes some space when it captures the changes that occurred on the volume. The changed blocks remain compressed, so snapshot capacity consumption is efficient.

Stagger snapshot schedules to ensure that application synchronization, I/O quiescing, database verification, and similar actions have time to complete before the next schedule starts.

All snap command options that accept a volume name as an argument will accept an optional --pool argument to differentiate those volumes that have the same name, but reside in different pools.

Synopsis

snap --help

snap --list
[--vol *vol_name*]
[--pool *pool_name*]
[--all]
[--unmanaged]
[--manually_managed]

snap --info *snap_name*
--vol *vol_name*
[--pool *pool_name*]
[--verbose]

snap --delete *snap_name*
--vol *vol_name*
[--pool *pool_name*]
[--force]
[--force_last_repl]

snap --edit *snap_name*
--vol *vol_name*
[--pool *pool_name*]
[--name *new_name*]
[--description *text*]

snap --online *snap_name*
--vol *vol_name*
[--pool *pool_name*]

snap --offline *snap_name*
--vol *vol_name*
[--pool *pool_name*]

[--force]

snap --cksum *snap_name*
--vol *vol_name*
[--pool *pool_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	[--vol *vol_name*]	List snapshots for the specified volume. Specify this suboption or --all, not both.	Guest
	[--pool *pool_name*]	List snapshots for the specified volume in the specified pool. The --pool is required if the volume name is not unique within the group.	
	[--all]	List snapshots for all volumes. Specify this suboption or --vol *vol_name*, not both.	
	[--unmanaged]	Limit list to unmanaged snapshots that will not be automatically deleted. Actions that lead to unmanaged snapshots include deletion of a protection schedule and dissociation of volume from volume collection. Manual snapshots are also considered unmanaged since they will not be deleted automatically.	
	[--manually_managed]	Limit list to manually managed snapshots. These snapshots will not be automatically deleted. Manually managed snapshots are either manually created, created by third party software, or created by the NimbleOS during restore and resize operations on a volume. You must use either the --all or --vol suboption to list manually managed snapshots.	
--info	*snap_name* --vol *vol_name*	Provide detailed information about the specified snapshot in the specified volume.	Guest
	[--pool *pool_name*]	List a specific snapshot from a set of snapshots for a specific volume. The --pool option is required if the volume name is not unique within the group.	
	--verbose	Give extra information about the specified snapshot, including metadata.	

Option	Arguments and Suboptions	Description	Role
--edit	*snap_name* --vol *vol_name*	Modify the specified snapshot. Specify at least one suboption.	Operator
	[--pool *pool_name*]	Modify the specified snapshot for the specified volume. The --pool is required if the volume name is not unique within the group.	
	[--name *new_name*]	Modify the name of the specified snapshot.	
	[--description *text*]	Modify the description of the specified snapshot or add a description to it. If there are spaces in the description, enclose the text in quotation marks.	
--online	*snap_name* --vol *vol_name*	Bring the specified snapshot online.	Operator
	[--pool *pool_name*]	Bring the specified snapshot online. The --pool is required if the volume name is not unique within the group.	
--offline	*snap_name* --vol *vol_name*	Take the specified snapshot offline. Initiators cannot connect to a snapshot that is offline. Taking a snapshot offline fails if there are open connections from initiators unless --force is used.	Operator
	[--pool *pool_name*]	Take the specified snapshot offline. The --pool is required if the volume name is not unique within the group.	
	[--force]	Forcibly take the specified snapshot offline even if it has open connections.	
--cksum	*snap_name* --vol *vol_name*	Compute the checksum of the specified snapshot to verify data integrity.	Operator
	[--pool *pool_name*]	Compute the checksum of the specified snapshot. The --pool is required if the volume name is not unique within the group.	
--delete	*snap_name* --vol *vol_name*	Delete the specified snapshot. Snapshots must be offline to be deleted.	Power User
	[--pool *pool_name*]	Delete the specified snapshot. The --pool is required if the volume name is not unique within the group.	
	[--force]	Forcibly delete a snapshot managed by an external agent.	
	[--force_last_repl]	Forcibly delete the last replicated snapshot.	

Examples

This example lists snapshots for the volume named *vdi*. The output is truncated to avoid line wrapping.

```
Nimble OS $ snap --list --vol vdi
---------+--------------------+----------+------+-------+------------
Volume    Snapshot             Size       Online Status  New Data
Name      Name                 (MiB)                      (MiB)
---------+--------------------+----------+------+-------+------------
vdi       vdi-volcoll-vdi-volcoll-2014-09-05::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-09-04::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-09-02::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-09-01::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-31::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-30::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-29::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-28::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-27::00:00:00.000    4194304...
vdi       vdi-volcoll-vdi-volcoll-2014-08-22::00:00:00.000    4194304...
```

This example provides detailed information about the specified snapshot.

```
Nimble OS $ snap --info 9lives-vc1-daily-2014-09-05::00:00:00.000 --vol 9lives
Name: 9lives-vc1-daily-2014-09-05::00:00:00.000
Volume name: 9lives
Pool: Pool2
Serial number: cdf437b74e248c706c9ce9003e217f89
iSCSI target:
Description: Replicated by protection policy 9lives-vc1 schedule daily
Size (MB): 994304
Status: Okay
Allow writes: No
Online: No
Offline reason: User
New data (MB): 0
New data compression: N/A
Number of connections: 0
iSCSI: 0
FC: 0
Schedule: daily
Origination group name/ID: 9lives
Is replica: Yes
Replication status: N/A
Created: Sep  5 2014 00:01:13
Last configuration change: Sep  5 2014 00:05:28
Access Control List:
Initiators connected:
```

This example modifies a snapshot to add a description and give it a simpler name.

```
Nimble OS $ snap --edit
    9lives-vc1-daily-2014-09-05::00:00:00.000 --vol 9lives
    --name snapforreplica05
    --description "snapshot taken on the 5th, for replication"
```

This example brings a snapshot online.

```
Nimble OS $ snap --online
    9lives-vc1-daily-2014-09-05::00:00:00.000 --vol 9lives
```

This example takes a snapshot offline.

```
Nimble OS $ snap --offline
    9lives-vc1-daily-2014-09-05::00:00:00.000 --vol 9lives
```

This example verifies data integrity of the snapshot named *9lives-vc1-daily-2014-09-05::00:00:00.000*.

```
Nimble OS $ snap --cksum 9lives-vc1-daily-2014-09-05::00:00:00.000
    --vol 9lives
INFO: Computing checksum for snapshot
9lives-vc1-daily-2014-09-05::00:00:00.000 of volume 9lives...
Checksum(v1.4): 0xffffffff
```

This example deletes the snapshot named *9lives-vc1-daily-2014-09-05::00:00:00.000*.

```
Nimble OS $ snap --delete 9lives-vc1-daily-2014-09-05::00:00:00.000
    --vol 9lives
```

snapcoll

The snapcoll command options manage snapshot collections. Snapshot collections are associated with volume collections.

Snapshot collections are replicated in the order that the collections were taken. If a snapshot cannot be replicated, the system indefinitely attempts to complete the replication. These pending snapshot collections are not deleted by the system, but they can be manually deleted after pausing the replication partner.

Except when initiated manually, snapshot collections are not created by the admin user; the system creates them after successful snapshot schedules complete. These snapshot collections are named with the volume collection name, schedule, and timestamp.

Downstream snapshot collections from synchronous replication can be identified by the .R appended to the end of the snapshot collection name.

Synopsis

snapcoll --help

snapcoll --list
[--volcoll *volcoll_name*]
[--unmanaged]

snapcoll --info *snapcoll_name*
[--volcoll *volcoll_name*]

snapcoll --delete *snapcoll_name*
[--volcoll *volcoll_name*]

snapcoll --edit *snapcoll_name*
[--volcoll *volcoll_name*]
[--name *new_name*]
[--description *text*]
[--replicate *{yes|no}*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--list	N/A	List all snapshot collections.	Guest
	[--volcoll *volcoll_name*]	List all snapshot collections associated with the specified volume collection. Use volcoll --list to get values for *volcoll_name*.	
	[--unmanaged]	Limit list to unmanaged snapshot collections that will not be automatically deleted. Deletion of a protection schedule leads to unmanaged snapshot collections. Manual snapshot collections are also considered unmanaged since they will not be deleted automatically.	
--info	*snapcoll_name*	Provide detailed information about the specified snapshot collection.	Guest
	[--volcoll *volcoll_name*]	Provide detailed information about the specified snapshot collection associated with the specified volume collection.	
--delete	*snapcoll_name*	Delete the specified snapshot collection.	Power User
	[--volcoll *volcoll_name*]	Delete the specified snapshot collection associated with the specified volume collection.	
--edit	*snapcoll_name*	Modify the specified snapshot collection.	Operator
	[--volcoll *volcoll_name*]	Modify the specified snapshot collection associated with the specified volume collection..	
	[--name *new_name*]	Modify the name of the specified snapshot collection.	
	[--description *text*]	Modify the plain-language description of the specified snapshot collection. This is useful to identify what the snapshot collection was designed to optimize. If there are spaces in the description, enclose the text in quotation marks.	
	[--replicate *{yes|no}*]	Enable or disable replication for the specified snapshot collection. **Note:** This command applies to snapshot replication only.	

Examples

This example lists the snapshot collections for the volume collection named *TestRep*.

```
Nimble OS $ snapcoll --list --volcoll TestRep
-----------------+-----------------------------------+-------+-----------
Volume Collection Snapshot Collection                Num        Replication
```

```
Name              Name                                    Snaps      Status
----------------+---------------------------------------+----------+-----------
TestRep           TestRep-FiveMinutes-2010-12-28::11:50:00.000   1 Pending
TestRep           TestRep-FiveMinutes-2010-12-28::11:50:00.000.R 1 Pending
TestRep          TestRep-FiveMinutes-2010-12-28::11:40:00.000   1 In-progress
TestRep           TestRep-FiveMinutes-2010-12-28::11:40:00.000.R 1 Complete
```

This example shows information about a collection of daily snapshots on a volume collection named *volume1--1754979654*.

```
Nimble OS $ snapcoll --info volume1--1754979654-daily-2014-07-16::00:00:00.000
--volcoll volume1--175497654
Name: volume1--175497654-daily-2014-07-16::00:00:00.000
Description: Created by protection policy volume1--175497654 schedule daily
Volume collection name: volume1--1754979654
Origination group name/ID: group-c20-array2
Is replica: No
Is complete: Yes
Is manual: No
Is externally triggered: No
Replication status: N/A
Replication started: N/A
Replication completed: N/A
Created: Jul 16 2014 00:00:00
Snapshots:
      Volume: volume1
              Snapshot: volume1--1754979654-daily-2014-07-16::00:00:00.000
Key/value pairs: None
Synchronous Replication Owner: default
```

This example changes the name and replication setting for a snapshot collection named *XCHG*.

```
Nimble OS $ snapcoll --edit XCHG --name EXCHANGE
--volcoll volume1--1754979654 --replicate no
```

This example deletes the snapshot collection named *TestRep* on the volume collection *Rep1*.

> **Note:** Snapshot collections cannot be deleted if they are in use by replication. If needed, pause or unconfigure replication to delete a snapshot collection.

```
Nimble OS $ snapcoll --delete TestRep --volcoll Rep1
```

software

The software command options manage the Nimble software that is installed on an array. You can download, upload, cancel an upload, and monitor software update status in real time.

You can have the currently installed version and one update version downloaded and ready for installation. Only one newer version can be stored at a time.

If your array does not have Internet access, you must download the software from the Nimble Storage Support site.

Synopsis

software --help

software --list

software --info *version*

[--view_license]

software --download *version*
[--force]

software --download_status
[--interval *seconds*]

software --cancel_download

software --precheck
[--force]

software --update
[--force]
[--accept_license]

software --update_status
[--verbose]

software --resume_update
[--accept_license]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all available software images.	Guest
--info	*version*	Provide detailed information about the specified software version.	Guest
	[--view_license]	Show the terms and conditions of use, the end user license agreement (EULA), and third-party software notices for the specified software version.	
--download	*version*	Start to download the specified software version.	Power User
	[--force]	Forcibly download the currently running software version.	
--download_status		Show the status of a software download.	Power User
	[--interval *seconds*]	Specify the number of seconds after which to refresh the status. Set this to 0 (zero) to disable the status refresh messages.	
--cancel_download	N/A	Cancel a software download.	Power User

Option	Arguments and Suboptions	Description	Role
--precheck		Perform verification checks on the array to verify that a software update can succeed.	Power User
	[--force]	Perform most verification checks, but skip the checks that can be overridden during the update by using the --force option then.	
--update		Start to update the array to the downloaded software version.	Power User
	[--force]	Forcibly start the software update, even if it might cause I/O errors on active volumes by restarting the active controller.	
	[--accept_license]	Accept the terms and conditions of use, the EULA, and third-party software notices.	
--update_status		Show the status of a software update.	Power User
	[--verbose]	Show the status of a software update on each array being updated.	
--resume_update		Resume a software update that timed out or failed before all member arrays were updated.	Power User
	[--accept_license]	Accept the terms and conditions of use, the EULA, and third-party software notices.	

Examples

This example lists available software versions on the array.

```
Nimble OS $ software --list
-----------------+-----------+-----------------------------------------------------
Version           Status      Size (MB)
-----------------+-----------+-----------------------------------------------------
2.1.5.0-118394-opt installed       808
2.1.5.0-117582-opt rollback        808
2.1.5.0-122684-opt downloaded  unknown
```

This example provides detailed information about software version *2.1.5.0-118394-opt*.

```
Nimble OS $ software --info 2.1.5.0-118394-opt
Version: 2.1.5.0-118394-opt
Status: installed
Size: 808 MB
Release Date: 2014-08-21
Release Notes: None
Fixes bugs:
```

This example forcibly downloads software version *2.1.5.0-118394-opt*, which is currently installed on the array.

```
Nimble OS $ software --download 2.1.5.0-118394-opt --force
```

This example shows the current status of a software download, refreshed every 10 seconds.

```
Nimble OS $ software --download_status --interval 10
```

This example cancels an in-progress software download.

```
Nimble OS $ software --cancel_download
Success
```

This example performs verification checks on a software download before installing the software version.

```
Nimble OS $ software --precheck
INFO: Software Update precheck passed.
```

This example begins a software update.

```
Nimble OS $ software --update
```

This example shows the status of an in-progress software update.

```
Nimble OS $ software --update_status
Updating group to version: 2.1.5.0-122684-opt
Group update start time Oct  5 2013 19:54:57
Group update end time N/A
     Updating array: array7
     Array update status: 1 of 7: Controller B is unpacking update package
```

sshkey

The sshkey command options manage secure shell (SSH) authorized keys. SSH uses public keys and cryptography to authenticate users.

When there is a public key on one end of a network connection and the matching private key on another, a user does not have to manually type the password. SSH keys make it possible to create scripts that perform certain tasks without having to manually log in to the array. You can use RSA and DSA algorithms. RSA is used both for encryption and signing. DSA is used only for signing.

Synopsis

sshkey --help

sshkey --list
[--user *user_name*]

sshkey --info *key_name*
[--user *user_name*]

sshkey --add *key_name*
--type *{dsa|rsa}*
--key *keystring*
[--user *user_name*]

sshkey --edit *key_name*
[--type *{dsa|rsa}*]
[--key *keystring*]
[--user *user_name*]

sshkey --delete *key_name*
[--user *user_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	[--user *user_name*]	List all currently stored SSH keys or the SSH key for the specified user. For each key, you will see the key name and type, but not the key string. To see the key string, use the --info option.	Administrator
--info	*key_name*	Show detailed information about the specified SSH key.	Administrator
	[*user_name*]	Show detailed information about the specified user's SSH key.	
--add	*key_name*	Specify the name of the SSH key. The key must exist before you add it.	Administrator
	--type *{dsa|rsa}*	Specify the type of SSH key authentication used.	
	--key *keystring*	Specify the DSA or RSA string that constitutes the SSH key. When you add the key, specify only the encrypted portion of the key. Do not include the ssh_dss or ssh_rsa at the start of the string or the *username@hostname* at the end.  **Important:** Do not include spaces at the start or end of the key. Include only the key string itself.	
	[--user *user_name*]	Add the SSH key for the specified user.	
--edit	*key_name*	Modify the specified SSH key. You can change the type (DSA or RSA) and the key string, but you cannot change the key name.	Administrator
	[--type *{dsa|rsa}*]	Modify the type of SSH key authentication used.	
	[--key *keystring*]	Modify the DSA or RSA string that constitutes the SSH key.	
	[--user *user_name*]	Modify the SSH key for the specified user.	
--delete	*key_name*	Delete the specified SSH key.	Administrator
	[--user *user_name*]	Delete the SSH key for the specified user.	

Examples

This example lists the current SSH keys.

```
Nimble OS $ sshkey --list
--------------------+------------------------------+---------
User                Name                           Type
--------------------+------------------------------+---------
cgrey               greybit                        rsa
guest               login                          rsa
```

This example lists only the guest user.

```
Nimble OS $ sshkey --list --user guest
--------------------+------------------------------+---------
User                Name                           Type
--------------------+------------------------------+---------
guest               login                          rsa
```

This example shows detailed information about the SSH key named *greybit*.

```
Nimble OS $ sshkey --info greybit
User: cgrey
Name: greybit
Type: rsa
Key: AAAAB3NzaC1yc2EAAAABIwAAAQEA1w+9D0j4w49/
xltKPbMKnwUEq7vlOntXx5vtaS3UGgeT+beCo2ERfRJCV
GIZZAUuA0COAwu3Y9J4wvvihuv7ri1UzR1U
```

This example adds an RSA SSH key to the system for the power user.

```
Nimble OS $ sshkey --add power --type rsa
--key AAAAB3NzaC1yc2EAAAABIwAAAQEA1w+9D0j4w49/
xltKPbMKnwUEq7vlOntXx5vtaS3UGgeT+beCo2ERfRJCVG
IZZAUuA0COAwu3Y9J4wvvihuv7ri1UzR1U --user poweruser
```

This example edits the key string for the RSA SSH key named *power*. The output is truncated to avoid line wrapping.

```
Nimble OS $ sshkey --edit power --type rsa
--key AAAAB3NzaC1yc2EAAAADAQABAAABAQDRt7osLN4w7CoibK58f9GYn/
l2JKqfS8Y0wy9BPeCGdOfNEv5TxSZBYNwMANZx2fC5yHzs7JJY85DsEce/65dUjRFDOhHoil...
uLyL8vfrxeVtPGjXWPPcKv64hTQRoyIxmW9Mb27nF2fJYy69WZNqvlJ/VAvjsnRuh/J42BVB...
+5gLzJ1j4GEnjBFm8IWnZ2PanhiAwIGZKwEXzU4iQfdujuPY5bOC... --user "power user"
```

This example deletes the RSA SSH key named *power*.

```
Nimble OS $ sshkey --delete power --user poweruser
```

stats

The stats command options let you view array statistics. You can specify the period of interest, items of interest, and display options. The default output includes a set of system utilization statistics.

All stats command options that accept a volume or folder name as an argument also accept an optional --pool argument to differentiate those volumes or folders that have the same name, but reside in different pools.

Synopsis

stats --help

stats
[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]
[--duration *[+]hh:mm[:ss]*]
[--array *array_name*]
[--disk {*all | 1,2,...*}]
[--shelf_location *shelf_location*]
[--net {*all | nic1,nic2,...*}]
[--fc {*all | fc1,fc2,...*}]
[--perf {*composite | vol*}]
[--latency]
[--iosize]
[--pool *pool*]
[--replication]
[--partner {*all | composite | partner1,partner2...*}]
[--vol {*all | composite | vol1,vol2,...*}]
[--pool *pool*]
[--folder {*all | composite | folder1,folder2,...*}]
[--pool *pool*]
[--interval *seconds*]
[--range]
[--show_time]
[--diff]
[--mb]
[--gb]
[--hdr *intervals*]
[--csv]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
	[--from *[[yyyy-]mm-dd,]hh:mm[:ss]*]	Specify the beginning time for which statistics are to be displayed. When used alone, the value defaults to the current time.	Guest
	[--to *[[yyyy-]mm-dd,]hh:mm[:ss]*]	Specify the ending time for which statistics are to be displayed. When used alone, the value defaults to the current time, creating a live statistics stream.	
	[--duration *[+]hh:mm[:ss]*]	Specify the length of time for which statistics are to be displayed. When used with --from, it specifies the duration starting at the --from time. When used with --to, it specifies the duration ending at the --to time. When used alone, it is equivalent to specifying --from "now minus duration", or if the specified duration begins with a '+', it is equivalent to specifying --to "now plus duration".	
	[--array *array_name*]	Display statistics from the specified array. For disk and net work statistics, this option is required when multiple arrays exist in the group. For volume statistics, this option is not required, and if not specified, group-wide aggregate values are displayed.	
	[--disk {*all* | *1,2,...*}]	Display statistics for disk utilization from the disk in the specified slot.	
	[--shelf_location *shelf_location*]	In the case of multiple disk shelves, use --shelf_location to see statistics for a particular disk in the specified shelf. If --shelf_location is not specified, disk statistics for all disks in the head shelf are displayed. For example, to see statistics for the disk in slot 7 of shelf 0 connected to controller A, you would enter stats --disk *7* --shelf_location *A.0*.	
	[--net {*all* | *nic1,nic2,...*}]	Display statistics for network utilization of all NICs or the specified NICs.	
	[--fc {*all* | *fc1,fc2,...*}]	Display statistics for all Fibre Channel ports or the specified ports.	
	[--perf {*composite* | *vol*}]	Display statistics for system or volume performance.	
	[--latency]	Display additional latency details with the --perf output.	
	[--iosize]	Display additional I/O details with the --perf output.	
	[--pool *pool*]		

Option	Arguments and Suboptions	Description	Role
		The --pool suboption is required if the volume name is not unique within the group. For example, to see statistics for vol1 in pool pool1, you would enter stats --perf *vol1* --pool *pool1*.	
	[--replication]	Display statistics for replications. Specify either the --partner suboption to specify one, all, or some replication partners, or the --vol suboption to specify one, all, or some volumes.	
	[--partner {*all* | *composite* | *partner1,partner2,...*}]	Display statistics for one, all (as a composite or individually), or some replication partners. Use partner --list to get *partner1,partner2,...* values.	
	[--vol {*all* | *composite* | *vol1,vol2,...*}]	Display statistics for one, all (as a composite or individually), or some volumes. Use vol --list to get *vol1,vol2,...* values.	
	[--pool *pool*]	Name of the pool where the volume or volumes reside.	
	[--folder {*all* | *composite* | *folder1,folder2,...*}]	Display statistics related to folders in a specific pool. For a single folder, specify the folder name. For a composite of all folders in a pool, specify the variable *composite*. For all folders, specify the variable *all*. Use folder --list to get *folder1,folder2,...* values.	
	[--pool *pool*]	The pool suboption is required to uniquely identify a folder within a group. For example, to see statistics for folder folder1 in pool pool1, you would enter stats --folder *folder1* --pool *pool1*.	
	[--interval *seconds*]	Display statistics in the specified intervals in seconds. If not specified, the latest values for the period are displayed.	
	[--range]	Display the minimum, maximum, and average values for the statistics.	
	[--show_time]	Display a timestamp for each interval.	
	[--diff]	Display statistics as differences from the previous interval.	
	[--mb]	Display disk and network byte transfers in mebibytes. The default is kibibytes.	
	[--gb]	Display disk and network byte transfers in gigabytes. The default is kibibytes.	
	[--hdr { *intervals* | *0*}]	Header lines are normally displayed every 20 intervals. Us this option to specify a different interval, or 0 to omit headers altogether.	

Option	Arguments and Suboptions	Description	Role
	[--csv]	Display statistics as a comma-separated list of values instead of as a space-separated list.	

Examples

> **Note:** To stop the scrolling display of statistics output when running the command interactively, press the Ctrl+C keys at any time.

These examples show stats command output, first without any parameters and then with defined parameters. The output is truncated to save space and prevent line wrapping.

```
Nimble OS $ stats
Stats from 2018-09-02,12:09:43 to 2038-01-18,19:14:07
       reads      writes    seqRead%   nsMemHit%  nsSsdHit%      netRxKB...
 12626306004 29069195568          25          17         39   7034286173...
 12626313650 29069208611          25          17         39   7034294229...
 12626315679 29069210890          25          17         39   7034297705...
 12626320142 29069216821          25          17         39   7034312508...
 12626324223 29069223926          25          17         39   7034316520...
 12626329577 29069231942          25          17         39   7034323314...
...
...


Nimble OS $ stats --interval 5
Will use per-second statistics
Stats from 2018-09-02,13:09:36 to 2038-01-18,19:14:07
       reads      writes    seqRead%   nsMemHit%  nsSsdHit%      netRxKB...
 12647799975 29095666635          25          17         39   7057545676...
 12647827863 29095708484          25          17         39   7057574592...
 12647853954 29095752416          25          17         39   7057603817...
...
...

Nimble OS $ stats --diff --mb
Stats from 2018-09-02,13:16:49 to 2038-01-18,19:14:07
       reads      writes    seqRead%   nsMemHit%  nsSsdHit%      netRxMB...
        6815        6308          39          15         25            4...
          --          --          --          --         --            8...
          --          --          --          --         --         1.17...
        6252        6592          36          17         27            4...
        6152        7729          37          14         30            4...
        5876        6888          38          14         31            9...
        8397       11263          36          11         37         0.41...
        2387        2827          39          16         33            6...
...
...

Nimble OS $ stats --perf greyhound --latency
Stats from 2018-09-02,13:21:51 to 2038-01-18,19:14:07
     io_type     0u-100u   100u-200u   200u-500u    500u-1m       1m-2m...
       reads          --          --          --         --          --...
     nsreads          --          --          --         --          --...
    seqreads          --          --          --         --          --...
      writes          --          --          --         --          --...
       reads        7941       13160        4613        671         433...
     nsreads        7144       12776        4591        661         429...
```

```
    seqreads         797         384          22         10           4...
      writes           0           0           0          0           0...
...
...
```

This example sets the start and end time between which to collect statistics. The output is truncated to prevent line wrapping.

```
Nimble OS $ stats --from 01:10 --to 01:20
Stats from 2018-09-02,01:10:00 to 2018-09-02,01:20:00
       reads      writes    seqRead%   nsMemHit%  nsSsdHit%      netRxKB...
 12655037664 29104182305          25          17         39   7085122631...
```

This example shows statistics for the eth3 NIC. The output is truncated to save space.

```
Nimble OS $ stats --net eth3
 Stats from 2018-08-12,09:00:28 to 2038-01-18,19:14:07
   dev            rxKB   rxPackets    rxErrors         txKB    txPackets...
  eth3       951076338   290812606           3     72776956    148602256...
  eth3       951076340   290812630           3     72776956    148602264...
  eth3       951076341   290812647           3     72776956    148602270...
  eth3       951076342   290812666           3     72776957    148602277...
  eth3       951076343   290812683           3     72776957    148602283...
...
...
```

This example shows statistics for the fc5 Fibre Channel port. The output is truncated to save space and prevent line wrapping.

```
Nimble OS $ stats --fc fc5
 Stats from 2018-08-12,09:02:30 to 2038-01-18,19:14:07
  dev          rxKB          txKB rxBBcreds txBBcreds Inv tx Word  CRC E...
  fc5  147154458457  135485861338        80         8         562       ...
  fc5  147154720749  135486151251        80         8         562       ...
  fc5  147154946494  135486425219        80         8         562       ...
...
...
```

This example shows statistics for the fc5 Fibre channel port in a comma-separated values (CSV) format showing gigabytes instead of the default of kilobytes. The output is truncated to save space and prevent line wrapping.

```
Nimble OS $ stats --fc fc5 --gb --hdr 30 --csv
Stats from 2018-08-13_11:32:34 to 2038-01-18_19:14:07
dev,rxGB,txGB,rxBBcreds,txBBcreds,Inv tx Word,CRC Errors,loss sync,loss ...
fc5,159457,148035,80,8,684,0,356,0,1
fc5,159457,148035,80,8,684,0,356,0,1
fc5,159457,148036,80,8,684,0,356,0,1
fc5,159458,148036,80,8,684,0,356,0,1
fc5,159458,148036,80,8,684,0,356,0,1
...
...
```

subnet

The subnet command options manage subnets.

Synopsis

subnet --help

subnet --list
[--netconfig *{active | backup | draft}*]

subnet --info *subnet_label*
[--netconfig *{active | backup | draft}*]

subnet --add *subnet_label*
--subnet_addr *network / netmask*
[--discovery_ipaddr *ipaddr*]
[--type *{mgmt | data | mgmt,data}*]
[--subtype *{iscsi | group | iscsi,group}*]
[--netzone_type *{evenodd | bisect | single}*]
[--netconfig *{active | backup | draft}*]
[--vlanid *id*]
[--mtu *mtu*]
[--failover *{enable | disable}*]

subnet --edit *subnet_label*
[--new_label *subnet_label*]
[--subnet_addr *network / netmask*]
[--discovery_ipaddr *ipaddr*]
[--type *{mgmt | data | mgmt,data}*]
[--subtype *{iscsi | group | iscsi,group}*]
[--netzone_type *{evenodd | bisect | single}*]
[--netconfig *{active | backup | draft}*]
[--vlanid *id*]
[--mtu *mtu*]
[--force_ip_update]
[--force_repl]
[--force_initiator_groups]
[--failover *{enable | disable}*]

subnet --remove *subnet_label*
[--netconfig *{active | backup | man subnetdraft}*]
[--force_repl]
[--force_ip_update]
[--force_unassign_nics]
[--force_initiator_groups]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all configured subnets.	Guest
	[--netconfig *{active | backup | draft}*]	List all configured subnets in the specified network configuration.	
--info	*subnet_label*	Provide detailed information about the specified subnet.	Guest
	[--netconfig *{active | backup | draft}*]	Provide detailed information about the specified subnet in the specified network configuration.	

Option	Arguments and Suboptions	Description	Role
--add	*subnet_label*	Add the specified subnet. The *subnet_label* can be up to 64 characters.	Power User
	--subnet_addr *network / netmask*	Specify the subnet address, expressed in Classless Inter-Domain Routing (CIDR) format.	
	[--discovery_ipaddr *ipaddr*]	Specify the subnet discovery IP address.	
	[--type *{mgmt | data | mgmt,data}*]	Specify whether the subnet is used for data, management, or data+management traffic. You can specify multiple types by using a comma-separated list.	
	[--subtype *{iscsi | group}*]	Specify the subnet subtype for a subnet on an iSCSI array. You can specify multiple subtypes by using a comma-separated list.	
	[--netzone_type *{evenodd | bisect | single}*]	Specify the Network Affinity Zone type for a subnet on an iSCSI array.	
	[--netconfig *name*]	Specify the network configuration to which to add the subnet.	
	[--vlanid *id*]	Specify the VLAN ID for the subnet. Valid values are in the 1-4094 range, inclusive.	
	[--mtu *mtu*]	Specify the maximum transmission unit (MTU) for the subnet. Valid values are in the 512-16000 range, inclusive.	
	[--failover *{enable | disable}*]	Specify whether to enable or disable controller failover on the subnet. By default, failover is enabled and the controller will failover when the standby controller has better connectivity to the subnet. Can be performed on a data subnet using the FC protocol.	

Option	Arguments and Suboptions	Description	Role
--edit	*subnet_label*	Modify the specified subnet. Specify at least one suboption.	Power User
	[--new_label *new_label*]	Modify the label for the specified subnet.	
	[--subnet_addr *network / netmask*]	Modify the subnet address, expressed in CIDR format.	
	[--discovery_ipaddr *ipaddr*]	Modify the subnet discovery IP address.	
	[--type *{mgmt | data | mgmt,data}*]	Modify the subnet type. You can specify multiple types by using a comma-separated list.	
	[--subtype *{iscsi | group}*]	Modify the subnet subtype for a subnet on an iSCSI array. You can specify multiple subtypes by using a comma-separated list.	
	[--netzone_type *{evenodd | bisect | single}*]	Modify the Network Affinity Zone type for a subnet on an iSCSI array.	
	[--netconfig *name*]	Modify the network configuration with which the subnet is associated.	
	[--vlanid *id*]	Modify the VLAN ID for the subnet. Valid values are in the 1-4094 range, inclusive.	
	[--mtu *mtu*]	Modify the MTU for the subnet. Valid values are in the 512-16000 range, inclusive.	
	[--force_ip_update]	Update the discovery IP address for the active network configuration.	
	[--force_repl]	Ignore inconsistent network configurations for replication partners when editing the specified subnet.	
	[--force_initiator_groups]	Ignore initiator groups assigned to subnets in the active network configuration when editing the specified subnet.	
	[--failover *{enable | disable}*]	Specify whether to enable or disable controller failover on the subnet. By default, failover is enabled and the controller will failover when the standby controller has better connectivity to the subnet. Can be performed on a data subnet using the FC protocol.	

Option	Arguments and Suboptions	Description	Role
--remove	*subnet_label*	Remove the specified subnet.	Power User
	[--netconfig *{active | backup | draft}*]	Specify the network configuration from which to remove the subnet.	
	[--force_repl]	Ignore inconsistent network configurations for replication partners when removing the specified subnet.	
	[--force_ip_update]	Remove the specified subnet and update the discovery IP address for the active network configuration.	
	[--force_unassign_nics]	Ignore the NIC assigned to the specified subnet in the active network configuration.	
	[--force_initiator_groups]	Ignore initiator groups assigned to subnets in the active network configuration when removing the specified subnet.	

Examples

This example lists configured subnets on a iSCSI array.

```
Nimble OS $ subnet --list
------------------+---------------+-------+-------------+----+-----
Label               Network         Type    Discovery IP   VLAN MTU
------------------+---------------+-------+-------------+----+-----
data1               192.0.2.0/23    Data    192.0.2.67     0    1500
mgmt-data           192.0.4.0/21    Mgmt    192.0.4.15     0    1500
Subnet-1.101.0.0    1.101.0.0/16    Data    N/A            2101 1500
Subnet-1.102.0.0    1.102.0.0/16    Data    N/A            2102 1500
Subnet-1.103.0.0    1.103.0.0/16    Data    N/A            2103 1500
Subnet-1.104.0.0    1.104.0.0/16    Data    N/A            2104 1500
Subnet-1.105.0.0    1.105.0.0/16    Data    N/A            2105 1500
Subnet-1.106.0.0    1.106.0.0/16    Data    N/A            2106 1500
Subnet-1.107.0.0    1.107.0.0/16    Data    N/A            2107 1500
Subnet-1.108.0.0    1.108.0.0/16    Data    N/A            2108 1500
Subnet-1.109.0.0    1.109.0.0/16    Data    N/A            2109 1500
Subnet-1.110.0.0    1.110.0.0/16    Data    N/A            2110 1500
```

This example provides detailed information about a configured subnet.

```
Nimble OS $ subnet --info data1
Label: data1
Network/Masklen: 192.0.2.0/23
Type: Data
Allow iSCSI: No
Allow group: Yes
Network Affinity Zone Type: None
Discovery IP: 192.0.2.67
VLAN Id: 0
MTU: 1500
```

This example adds a subnet to the *draft* network configuration.

```
Nimble OS $ subnet --add Subnet-1.102.0.0
--subnet_addr 1.102.0.0/16 --type data --netconfig draft
--vlanid 2101 --mtu 1500
```

This example modifies a previously added subnet in the *draft* network configuration.

```
Nimble OS $ subnet --edit Subnet-1.102.0.0
--netconfig draft --mtu 9000
```

This example removes a subnet and ignores the initiator groups assigned to it.

```
Nimble OS $ subnet --remove Subnet-1.102.0.0  --force_initiator_groups
```

support

The support command allows you to manage temporary passwords for support access.

Synopsis

support --help

support
[--info]
[--show_ciphertext *{array_name | SN}*]
[--cycle]

Options

Table

Option	Arguments and Suboptions	Description	Role
[--help]	N/A	Show usage information for the command.	Admin
[--info]	N/A	Show enabled parameters.	Admin
[--show_ciphertext]	{array_name | SN}	Show array information and ciphertext blobs for all secure accounts on the defined array.	Admin
[--cycle]		Cycle the support passwords for all support accounts on all arrays in the group.	Admin

Examples

These examples show the support command output. The output may be truncated to save space and prevent line wrapping.

This example show the read secure password mode:

```
Nimble OS $ support --info
Password mode: ciphertext
```

This example shows the ciphertext successfully displayed:

```
Nimble OS $ support --show_ciphertext sbillimor-vma-2
Name: c20-array1
Serial: AC-100159
Model: CS210
User: root
Blob:
--- Begin tpd blob ---
0500
898410055e2c5c33b584b1a80e082f333d2de6e2cb...
 --- End tpd blob ---
User: nsupport
Blob:
--- Begin tpd blob ---
0500
3132141f9033951a903513b153156bef9351df8351...
 --- End tpd blob ---
```

timezone

The timezone command options manage time zone settings.

Synopsis

timezone --help

timezone --list

timezone --set *zone_name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List all valid time zones for the Nimble array.	Guest
--set	*zone_name*	Set the specified time zone on an array. Use timezone --list to get a list of valid *zone_name* values.	Power User

Examples

This example lists the valid time zones. The output is truncated to save space.

```
Nimble OS $ timezone --list
Asia/Baghdad
Asia/Bahrain
Asia/Baku
Asia/Bangkok
Asia/Beirut
Asia/Bishkek
Asia/Brunei
Asia/Calcutta
```

```
Asia/Choibalsan
...
...
Pacific/Samoa
Pacific/Tahiti
Pacific/Tarawa
Pacific/Tongatap
Pacific/Truk
Pacific/Wake
Pacific/Wallis
Pacific/Yap
```

This example sets the time zone on the array to the time in Stockholm, Sweden.

```
Nimble OS $ timezone --set Europe / Stockholm
```

useradmin

The useradmin command options manage user accounts.

Synopsis

useradmin --help

useradmin --list
[--role *{administrator | poweruser | operator | guest}*]

useradmin --info *user_name*

useradmin --add *user_name*
[--description *text*]
[--full_name *text*]
[--email_addr *email_address*]
[--role *text*]
[--inactivity_timeout *minutes*]

useradmin --edit *user_name*
[--description *text*]
[--full_name *text*]
[--email_addr *email_address*]
[--role *{administrator | poweruser | operator | guest}*]
[--inactivity_timeout *minutes*]

useradmin --disable *user_name*

useradmin --enable *user_name*

useradmin --remove *user_name*

useradmin --passwd
[--user *user_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	[--role *{administrator | poweruser | operator | guest}*]	List either all user accounts or user accounts that are assigned to the specified role.	Administrator
--info	*user_name*	Show either detailed information about the logged-in user if you specify your own *user_name* or any specified user if you are logged in as a user who is assigned to the Administrator role.	Guest
add	*user_name*	Add an account for the specified user. The *user_name* must be alphanumeric and start with a letter, be 1 to 32 characters, and contain no spaces.	Administrator
	[--description *text*]	Specify a plain-text description of the user account. The *text* can be blank or can contain up to 255 characters. If there are spaces in the description, enclose the text in quotation marks.	
	[--full_name *text*]	Specify the full name of the user account. The *text* can be blank or start with a letter, contain up to 64 alphanumeric characters, dashes, underscores, or apostrophes, no spaces, no periods.	
	[--email_addr *email_address*]	Specify an email address for the user.	
	[--role *{administrator|poweruser|operator|guest}*]	Specify the role the user has.	
	[--inactivity_timeout *minutes*]	Set the number of minutes that must pass without any user activity before a timeout occurs and this user is logged out. If you do not specify this option, the default of 30 minutes is used. The *minutes* cannot exceed the inactivity timeout set for the group.	

Option	Arguments and Suboptions	Description	Role
--edit	*user_name*	Modify information about the logged-in user if you specify your own *user_name* or any specified user if you are logged in as a user who is assigned to the Administrator role. Specify at least one suboption.	Guest
	[--description *text*]	Modify the description of the specified user account. If there are spaces in the description, enclose the text in quotation marks.	
	[--full_name *text*]	Modify the full name of the user account.	
	[--email_addr *email_address*]	Modify the email address of the user account.	
	[--inactivity_timeout *minutes*]	Modify the number of minutes that must pass without any user activity before a timeout occurs and this user is logged out.	
	[--role *{administrator | poweruser | operator | guest}*]	Modify the role the user has. Logged-in users cannot change their own role. You must be assigned to the Administrator role to modify a user's role.	Administrator
--disable	*user_name*	Disable the specified user account.	Administrator
--enable	*user_name*	Enable the specified user account. This option succeeds even if the account is currently enabled.	Administrator
--remove	*user_name*	Remove the specified user account.	Administrator
--passwd	[--user *user_name*]	Modify either your own password if you do not specify the optional --user option or the password of the specified user if you are logged in as a user who is assigned to the Administrator role and you do specify --user. The new password value is case sensitive, must be at least eight characters, must be different from the current password, and cannot include left or right square brackets ([ ]), ampersands (&), semicolons (;), or back single quotation mark (`).	Guest

Examples

This example lists the user accounts on anuseradmin array.

```
Nimble OS $ useradmin --list
------------------+---------------+--------+---------+--------------------
Name                Role            Status   Logged in Last login
------------------+---------------+--------+---------+--------------------
```

```
admin               administrator   Enabled  Y         Sep  2 2014 13:47:01
chris               operator        Enabled  N         Aug  1 2014 14:45:57
guest               guest           Enabled  N         Aug 22 2014 10:25:45
operator            operator        Enabled  N         Aug 24 2014 13:53:23
pAdmin              administrator   Enabled  Y         Sep  2 2014 14:36:53
pGuest              guest           Enabled  N         Aug  7 2014 12:28:50
pOper               operator        Enabled  N         Aug  4 2014 15:07:39
pPower              poweruser       Enabled  N         Sep  2 2014 09:36:46
vcplugin            poweruser       Enabled  N         Aug 22 2014 10:01:26
```

This example shows the current settings for the *admin* user.

```
Nimble OS $ useradmin --info admin
Name: admin
Role: administrator
Status: enabled
Logged in: true
Description: Administrator
Full name: Administrator
Email address:
Inactivity timeout: 30 minute(s)
Last login: Apr 13 2015 09:33:42
Last logout: Apr 13 2015 10:13:59
Creation time: N/A
Last modified: 1
Id: 1
```

This example modifies the description and inactivity timeout for the user account named *chris*.

```
Nimble OS $ useradmin --edit chris --description "Chris Green" --inactivity_timeout
 90
```

This example modifies the password of the logged in user.

```
Nimble OS $ useradmin --passwd
Enter current password:
Enter new password:
Retype new password:
Nimble OS $
```

userauth

The userauth command manages Active Directory groups, users and domains.

Synopsis

userauth --help

userauth --list

userauth --info *domain_name*

userauth --join *domain_name*
--domain_user *domain_user_name*
[--ou *ou_name*]
[--computer_name *computer_name*]
[--netbios_name *netbios_name*

userauth --leave *domain_name*
--domain_user *domain_user_name*
[--force]

userauth --enable *domain_name*

userauth --disable *domain_name*

userauth --add_group *group_name*
--domain *domain_name*
[--role *role*]
[--description *description*]
[--inactivity_timeout *minutes*]

userauth --remove_group *group_name*
--domain *domain_name*

userauth --edit_group *group_name*
--domain *domain_name*
[--role *role*]
[--description *description*]
[--inactivity_timeout *minutes*]

userauth --enable_group *group_name*
--domain *domain_name*

userauth --disable_group *group_name*
--domain *domain_name*

userauth --list_group
--domain *domain_name*

userauth --info_group *group_name*
--domain *domain_name*

userauth --test_group *group_name*
--domain *domain_name*

userauth --test_user *user_name*
--domain *domain_name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Administrator
--list		List the Active Directory domains.	Administrator
--info	*domain_name*	Show detailed information about the *domain_name* specified.	Administrator

Option	Arguments and Suboptions	Description	Role
--join	*domain_name*	Join the specified Active Directory domain.	Administrator
	[--domain_user *domain_user_name*]	Name of the Active Directory user with privileges to join the Active Directory domain. This is a mandatory argument.	
	[--ou *ou_name*]	Organization Unit, the location for the computer account on the Active Directory server. If not specified, the default location is Computers.	
	[--computer_name *computer_name*]	The name of the computer account in the domain controller on the Active Directory server. If not specified, the default value is the last 15 characters of the group name of the array.	
	[--netbios_name *netbios_name*]	The netbios name for the Active Directory domain. The default is the short name of the domain name.	
--leave	*domain_name*	Leave the specified Active Directory domain.	Administrator
	[--domain_user *domain_user_name*]	The name of the Active Directory user with privileges to join the Active Directory domain. This is a mandatory argument.	
	[--force]	Forcibly leave an Active Directory domain.	
--enable	*domain_name*	Enable the specified Active Directory domain.	Administrator
--disable	*domain_name*	Disable the specified Active Directory domain.	Administrator
--add_group	--domain*domain_name*	Add a group to the specified Active Directory domain.	Administrator
	[--role *role*]	The array role to assign to the users that belong to the Active Directory group. If not specified, the default role is guest. The options are: {administrator|poweruser|operator|guest}.	
	[--description *description*]	The description of the Active Directory group.	
	[--inactivity_timeout *minutes*]	The amount of time that the user session is inactive before timing out in minutes. The inactivity timeout for Active Directory groups cannot be more than the group user inactivity time-out for the array.	
--remove_group	--domain*domain_name*	Remove a group from the specified Active Directory domain.	Administrator

Option	Arguments and Suboptions	Description	Role
--edit_group	--domain*domain_name*	Edit a group on the specified Active Directory domain.	Administrator
	[--role*role*]	Modify the role of the Active Directory group.	
	[--description*description*]	Modify the description of the Active Directory group.	
	[--inactivity_timeout*minutes*]	Modify the user inactivity timeout for the Active Directory group.	
--enable_group	--domain*domain_name*	Enable a group on the specified Active Directory domain. This is a mandatory argument.	Administrator
--disable_group	--domain*domain_name*	Disable a group on the specified Active Directory domain. This is a mandatory argument.	Administrator
--list_group	--domain*domain_name*	List the groups on the specified Active Directory domain. This is a mandatory argument.	Administrator
--info_group	--domain*domain_name*	List detailed information about a specified Active Directory domain. This is a mandatory argument.	Administrator
--test_group	--domain*domain_name*	The Active Directory domain that the group belongs to. This is a mandatory argument.	Administrator
--test_user	--domain*domain_name*	The Active Directory domain that the user belongs to. This is a mandatory argument.	Administrator

Examples

This example lists the user accounts on an array.

```
Nimble OS $ userauth --list
------------------+-----------------+---------+-
Name                Connection Status   Enabled
------------------+-----------------+---------+-
win2k8r2ad.net  Connected   Yes
```

This example shows the current settings for the configured Active Directory Domains.

```
Nimble OS $ userauth --info win2k8r2ad.net
Active Directory Domain Name: win2k8r2ad.net
Organizational Unit: AD default location
Computer Name: esc-sjc-array343
Netbios Name: win2k8r2ad
Status: Successfully connected to Domain Controller.
```

This example lists the groups on the specified Active Directory domain.

```
Nimble OS $ userauth --list_group --domain win2k8r2ad.net
-------------------------+--------------+---------+--------------------
```

```
Name                              Role        Enabled   Description
--------------------------------+---------------+---------+----------------------
eng-qa-security-admin          administrator    Yes
```

userpolicy

The userpolicy command enables administrators to set global security policies for a group.

Synopsis

userpolicy --help

userpolicy --edit --allowed_attempts*num_allowed_attempts*
[--min_length *minimum_length*]
[--upper *num_uppercase_chars*]
[--lower *num_lowercase_chars*]
[--digit *num_digits*]
[--special *num_special_chars*]
[--previous_diff *num_chars_change*]
[--no_reuse *num_no_reuse_last*]
[--max_sessions *num_max_group_sessions*]

userpolicy --info

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Provides program help.	Guest
--edit	--allowed_attempts *num_allowed_attempts*	Number of failed authentication attempts allowed before your account is locked.	Administrator
	[--min_length *minimum_length*]	Minimum number of characters required for a valid password.	
	[--upper *num_uppercase_chars*]	Minimum number of uppercase characters required for a valid password.	
	[--lower *num_lowercase_char*]	Minimum number of lowercase characters required for a valid password.	
	[--digit *num_digits*]	Minimum number of numerical characters required for a valid password.	
	[--special *num_special_chars*]	Minimum number of special characters required for a valid password.	
	[--previous_diff *num_chars_change*]	Minimum number of characters that must be different from the previous password.	
	[--no_reuse *num_no_reuse_last*]	The number of times that a password must change before you can reuse an old password.	
	[--max_sessions *num_max_group_sessions*]	Maximum number of user sessions for a group.	

Option	Arguments and Suboptions	Description	Role
--info	N/A	Displays information about the configured security policy.	Administrator

usersession

The usersession command options manage user sessions.

Synopsis

usersession --help

usersession --list

usersession --info
[--id *session_id*]

usersession --kill *session_id*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List either the current user session if you are logged in as a user with any role other than Administrator or all user sessions if you are logged in as a user who is assigned to the Administrator role.	Guest
--info	[--id *session_id*]	Show either detailed information about the current session for the logged-in user if you do not specify the optional --id option or the specified user session if you do specify --id. If you are logged in as a user who is assigned to the Administrator role, you can specify something other than the current session by including the --id *session_id* option.	Guest
--kill	*session_id*	End the specified user session. The user is not logged out, but cannot take any more actions in that session. You cannot end the current user session. If you try to, you see a message saying that the Session cannot commit suicide.	Administrator

Examples

This example lists current user sessions. In this example, only the admin user is logged in to run the command. The admin user is also logged in from multiple separate sessions. The session ID (SID) is a hash (digest) rather than the real SID. The Source IP Addr is the IP address of the computer from which the user logged in to the command line interface (CLI), the

graphical user interface (GUI), or both. A single user can have multiple user sessions. The Expiry Time (secs) is the number of seconds until this session ends if no user action is taken. Any user action resets the clock.

```
Nimble OS $ usersession--list
---------------+----------------------------+---------------+--------------
User            Session ID                   Source IP Addr. Expires in
                                                             (secs)
---------------+----------------------------+---------------+--------------
admin           6worApuItdHDcC39dJOnTULgfp0= 192.0.2.177     127

admin           9ukXV/OjUA0T/0mIcOxiZ6J53WM= 192.0.2.177     201

admin           Yl9zye/rXZS7z9wSmBWBZmgoimE= 198.51.100.18   727
```

This example shows details about the current session for the logged-in user.

```
Nimble OS $ usersession --info
Session ID: Yl9zye/rXZS7z9wSmBWBZmgoimE=
User name: admin
Source IP: 198.51.100.18
Expiry Time (secs): 727
```

This example shows details about the specified session. The admin user runs this version of the command.

```
Nimble OS $ usersession --info --id 6worApuItdHDcC39dJOnTULgfp0=
Session ID: 6worApuItdHDcC39dJOnTULgfp0=
User name: admin
Source IP: 192.0.2.177
Expiry Time (secs): 127
```

This example kills the GUI session with a SID of *9ukXV/OjUA0T/0mIcOxiZ6J53WM=* for the admin user. Even though this example is somewhat contrived, an administrator user can kill another admin session that is not the command-line interface from which the kill action is initiated.

```
Nimble OS $ usersession--kill 9ukXV/OjUA0T/0mIcOxiZ6J53WM=
```

vcenter

Synopsis

The vcenter command options allow you to manage VMware extensions, and to add, edit, and delete the vCenter appserver.

vcenter --help

vcenter --list

vcenter --info *vcenter_name*

vcenter --register *vcenter_name*
[--extension {*web* | *vasa*}]

vcenter --unregister *vcenter_name*
[--extension {*web* | *vasa*}]

vcenter--status *vcenter_name*
[--extension {*web* | *vasa*}]

vcenter --add *name*
[--hostname {*hostname* | *ipaddr*}]
[--port_number *port_number*]
[--username *user_name*]
[--password *password*]
[--description *description*]
[--subnet_label*subnet_label*]

vcenter--edit *vcenter_name*
[--name*vCenter_name*]
[--username*user_name*]
[--password*password*]
[--description*description*]

vcenter--remove*vcenter_name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	N/A	List detailed information about a specified VMware vCenter, and the status of Nimble Storage vCenter extensions in the vCenter.	Guest
--info	*vcenter_name*	VMware vCenter and Nimble Storage's vCenter extensions status information based on vCenter name.	Guest
--register	*vcenter_name*	Register the extension with vCenter.	Administrator
	[--extension {*web* | *vasa*}]	Extension to register the vCenter.	
--unregister	*vcenter_name*	Unregister the extension with vCenter.	Administrator
	[--extension {*web* | *vasa*}]	Extension to unregister the vCenter.	
--status	*vcenter_name*	Status of an extension from the vCenter.	Operator
	[--extension {*web* | *vasa*}]	Extension to get the status.	
--add		Add VMware vCenter for this group.	Administrator
	--name	VMware vCenter name.	
	[--hostname {*hostname* | *ipaddr*}]	VMware vCenter server hostname or IP address.	
	[--port_number *port_number*]	VMware vCenter port number.	
	[--username *user_name*]	VMware vCenter username.	
	[--password *password*]	VMware vCenter user password.	
	[--description *description*]	VMware vCenter description.	
	[--subnet_label *subnet_label*]	Subnet label used for registration.	

Option	Arguments and Suboptions	Description	Role
--edit	*vcenter_name*	Edit vCenter attributes.	Administrator
	[--name *vcenter_name*]	VMware vCenter name to be edited.	
	[--username *user_name*]	Modify the VMware vCenter username.	
	[--password *password*]	Modify the VMware vCenter user password.	
	[--description *description*]	Modify the VMware vCenter description.	
--remove	*vcenter_name*	Remove specified vCenter server from this group.	Administrator

Examples

This example changes the vcenter server password.

```
Nimble OS $ vcenter --edit vserver1 --username testuser --password testpass
```

This example lists all the vcenters on the array.

```
Nimble OS $ vcenter --list
---------------+-------------+------+---------------------+----------+----------------------+---------
Name                 Host Name     Port    UserName                        Description
 Registered Extensions    Subnet Label
---------------+-------------+------+---------------------+----------+----------------------+---------
10.18.238.201   10.18.238.201 443    admin@vsphere.local   [web, vasa]
    mgmt-data
10.18.238.248   10.18.238.248 443    admin@vsphere.local   [web]
    mgmt-data
10.18.238.225   10.18.238.225 443    admin@vsphere.local   [web, vasa]
    mgmt-data
```

This example lists all the information about the specified vcenter, 10.18.238.201.

```
Nimble OS $  vcenter --info 10.18.238.201
------------+----------+----+-----------------+--------+----------+---------+--------+-----------
Name                 Host Name     Port    UserName                        Description
 Subnet Label Web Client  Vasa Provider
------------+----------+----+-----------------+--------+----------+---------+--------+-----------
10.18.238.201  10.18.238.201 443  admin@vsphere.local
    mgmt-data    Registered  Registered
```

vm

The vm command options allow you to manage VMware VVOL-based virtual machines.

Synopsis

VM objects are added using VMware's New Virtual Machine wizard, but can be managed using NimbleOS.

vm--help
vm--list [--deleted] [--clone_parents]vm--info {*name* | *id*}
vm--restore {*name* | *id*}
vm--destroy {*name* | *id*}

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list		List VMware virtual machines.	
	--deleted	Show only deleted virtual machines.	
	--clone_parents	Include VMs that are clone parents.	
--info	{name | id}		
--restore	{name | id}	Restore a deleted virtual machine.	
--destroy	{name | id}	Destroy a deleted virtual machine. This will also delete the volumes that comprise this VM.	

This example lists all of the VMs on the array.

```
Nimble OS $ vm --list
----------------------------------------+------+-----------+--------------------
Name                                      Type  Deleted     Time Until
Cleanup
----------------------------------------+------+-----------+--------------------
RajVvolTest                              VVol        No           N/A
proVm1                                   VVol        No           N/A
testVm3                                   VVol        No           N/A
vm1                                         VVol        No           N/A
```

This example returns information about the specified VM, vm1.

```
Nimble OS $ vm --info vm1
VM ID: 502570ce-7e48-d602-b1ee-37f1ab66c428
VM name: vm1
Volumes:
             vm1
              vm1.vmdk
```

version

The version command options display the version of the NimbleOS (NOS) software that is installed on an array or a group of arrays.

Synopsis

version --help

version
[--array *name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
[--array]	*name*	Specify the array for which to show version information. Use array--list to get valid *name* values. If you do not specify this option, then the version information is for the group.	Guest

Examples

This example displays the current version of the NOS software on a group of arrays.

```
Nimble OS $ version
INFO: Version information is for the group.
2.1.4.0-100755-opt
```

This example displays the current version of the NOS software on an array named *greyhound*.

```
Nimble OS $ version --array greyhound
INFO: Version information is for array greyhound
3.1.4.0-100755-opt
```

vmwplugin

The vmwplugin command options manage the NimbleOS plugin for the VMware vCenter.

Use the plugin to manage datacenter and datastore array-specific functions, such as cloning, creating, and managing datastores that reside on HPE Nimble Storage arrays. To install and use the VMware plugin, provide vCenter host and authentication information.

> **Note:** The plugin is not supported for multiple datastores that are located on one LUN, for one datastore that spans multiple LUNs, or if the LUN is located on a non-HPE Nimble Storage array. For these conditions, use the **vcenter** command.

Synopsis

vmwplugin --help

vmwplugin --list
--username *user_name*
--password *password*
--server *server*
[--port_number *port_number*]

vmwplugin --register
--username *user_name*
--password *password*
--server *server*
[--subnet_label *subnet_label*]
[--port_number *port_number*]
[--client {*web|thick*}]

vmwplugin --unregister
--username *user_name*
--password *password*
--server *server*
[--port_number *port_number*]
[--client {*web*|*thick*}]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	The --username, --password, and --server suboptions are mandatory.	List the VMware vCenter plugins.	Guest
	--username *user_name*	Specify the username of an account that can access the plugin.	
	--password *password*	Specify the VMware vCenter password for the specified username.	
	--server *server*	Specify the VMware vCenter fully qualified vCenter server name or IP address.	
	[--port_number *port_number*]	Specify the VMware vCenter port number. If you do not specify this option, the default port number 443 is used.	
--register	The --username, --password, and --server suboptions are mandatory.	Register a VMware vCenter plugin.	Power User
	--username *user_name*	Specify the username of an account that can access the plugin.	
	--password *password*	Specify the VMware vCenter password for the specified username.	
	--server *server*	Specify the VMware vCenter fully qualified vCenter server name or IP address.	
	[--subnetlabel *subnet_label*]	Specify a subnet label to use for registration. Example: *mgmt-data*. If you do not specify this option, no subnet label is used.	
	[--port_number *port_number*]	Specify the VMware vCenter port number. If you do not specify this option, the default port number 443 is used.	
	[--client {*web*|*thick*}]	Specify whether to allow access to the VMware vCenter plugin through the web-based or thick (vSphere) client.	

Option	Arguments and Suboptions	Description	Role
--unregister	The --username, --password, and --server suboptions are required.	Unregister a VMware vCenter plugin.	Power User
	--username *user_name*	Specify the username of an account that can access the plugin.	
	--password *password*	Specify the VMware vCenter password for the specified username.	
	--server *server*	Specify the VMware vCenter fully qualified vCenter server name or IP address.	
	[--port_number *port_number*]	Specify the VMware vCenter port number. If you do not specify this option, the default port number 443 is used.	
	[--client {*web*|*thick*}]	Specify whether to remove access to the VMware vCenter plugin through the web-based or thick (vSphere) client.	

Examples

This example lists configured VMware vCenter plugins.

```
Nimble OS $ vmwplugin --list --username vcenteradmin --password mypassword
--server 203.0.113.1
```

This example registers the VMware vCenter plugin. Because the --port option is not specified, the default of 443 is used.

```
Nimble OS $ vmwplugin --register --username vcenteradmin
--password mypassword --server 203.0.113.1
```

This example unregisters a VMware vCenter plugin.

```
Nimble OS $ vmwplugin --unregister --username vcenteradmin
--password mypassword --server 203.0.113.1
```

vol

The vol command options manage volumes.

Volumes are the basic storage units from which the total capacity of an array is apportioned. The array is sectioned into volumes. The number of volumes per array depends on storage allocation. In the context of the HPE Nimble Storage array, the term volume and LU (logical unit) are synonymous and interchangeable. A number identifies each LU, so you might see the term *LUN* used, too.

All vol commands that accept a volume name as an argument will accept an optional --pool argument to differentiate those volumes that have the same name but reside in different pools.

Note:

Keep the following points in mind when you use the encryption feature:

- When you create a volume, either the group configuration or your volume specifications determine whether the volume is encrypted. After volume creation, encryption on that volume cannot be changed.
- If you clone a volume, the clone inherits the encryption status of the parent volume. You cannot change that initial encryption status on the clone.

- If you lose the passphrase for an encrypted volume, the volume is inaccessible and its data is irretrievable.
- Performance on CS2xx arrays can be slow when accessing encrypted volumes.

Synopsis

vol --help

vol --list
[--pool *pool_name*]
[--folder*folder_name*]
[--initiatorgrp *igrp_name*]
[--open_access]
[--agent_type {*none*|*smis*|*vvol*|*openstack*|*openstackv2*}]
[--moving]
[--cache_pinned {*yes*|*no*}]
[--deduped{*yes*|*no*}]

vol --info *name*
[--pool *pool_name*]
[--verbose]

vol --create *name*
--size *mebibytes*
[--description *text*]
[--perfpolicy *name*]
[--cache_pinned {*yes*|*no*}]
[--thinly_provisioned {*yes*|*no*}]
[--limit *percent*]
[--start_offline]
[--apply_acl_to {*volume*|*snapshot*|*both*}]
[--chapuser *user_name*]
[--initiatorgrp *group_name*]
[--lun *lun*]
[--multi_initiator {*yes*|*no*}]
[--pool *pool_name*]
[--folder *folder_name*]
[--agent_type {*none*|*smis*}]
[--encryption_cipher {*aes-256-xts*|*none*}]
[--dedupe_enabled {*yes*|*no*}]
[--iscsi_target_scope {*volume*|*group*}]
[--iops_limit *iops*]
[--mbps_limit *mbps*]

vol --delete *vol_name*
[--pool *pool_name*]
[--force]

vol --edit *vol_name*
[--name *new_name*]
[--description *text*]
[--size *mebibytes*]
[--perfpolicy *name*]
[--cache_pinned {*yes*|*no*}]
[--readonly {*yes*|*no*}]
[--force]

[--thinly_provisioned {*yes*|*no*}]
[--limit *percent*]
[--multi_initiator {*yes*|*no*}]
[--agent_type { *none*|*smis*}]
[--dedupe_enabled {*yes*|*no*}]
[--iscsi_target_scope {*volume*|*group*}]
[--iops_limit *iops*]
[--mbps_limit *mbps*]
[--pool *pool_name*]

vol --online *vol_name*
[--pool *pool_name*]

vol --offline *vol_name*
[--pool *pool_name*]
[--force]

vol --snap *vol_name*
--snapname *name*
[--description *text*]
[--start_online]
[--allow_writes]
[--pool *pool_name*]

vol --restore *vol_name*
--snapname *snap_name*
[--pool *pool_name*]

vol --clone *vol_name*
--snapname *snap_name*
--clonename *clone_name*
[--folder *folder_name*]
[--description *text*]
[--readonly {*yes*|*no*}]
[--thinly_provisioned {*yes*|*no*}]
[--limit *percent*]
[--start_offline]
[--apply_acl_to {*volume*|*snapshot*|*both*}]
[--chapuser *user_name*]
[--initiatorgrp *group_name*]
[--lun *lun*]
[--multi_initiator {*yes*|*no*}]
[--cache_pinned {*yes*|*no*}]
[--dedupe_enabled {*yes*|*no*}]
[--iops_limit *iops*]
[--mbps_limit *mbps*]
[--pool *pool_name*]

vol --assoc *vol_name*
--volcoll *name*
[--pool *pool_name*]

vol --dissoc *vol_name*
[--pool *pool_name*]
[--force]

vol --addacl *vol_name*
--apply_acl_to {*volume|snapshot|both*}
[--chapuser *user_name*]
[--initiatorgrp *group_name*]
[--lun *lun*]
[--pool *pool_name*]

vol --removeacl *vol_name*
[--apply_acl_to {*volume|snapshot|both*}]
[--chapuser *user_name*]
[--initiatorgrp *group_name*]
[--pool *pool_name*]

vol --claim *vol_name*
[--partner *partner_name*]
[--pool *pool_name*]

vol --move *vollist*
--dest_pool *pool_name*
[--dest_folder *folder_name*]
[--pool *pool_list*]

vol --abort_move *vol_name*
[--pool *pool_name*]

vol --assign_folder *vol_name*
--folder *folder_name*
[--pool *pool_name*]

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest

Option	Arguments and Suboptions	Description	Role
--list		List the volumes on the array where you run the command.	Guest
	[--pool *pool_name*]	Limit the list of volumes to a specific storage pool.	
	[--folder *folder_name*	Limit the list of volumes to folders with the specified name. If a pool is not specified, the folder name can refer to multiple folders (one in each pool). If a pool is specified, the pool and folder names together refer to a single folder.	
	[--initiatorgrp *igrp_name*]	Limit the list of volumes to the specified initiator group.	
	[--open_access]	Limit the list of volumes to those that have open access.	
	[--agent_type {*none*|*smis*|*vvol*|*openstack*|*openstackv2*}]	Limit the list of volumes to the specified external management agent type. The *smis* agent type lists volumes on Storage Management Initiative Specification (SMI-S) enabled hosts. The *vvol* agent type lists virtual volumes. The *openstack* agent type lists volumes on Openstack enabled hosts. The *openstackv2* agent type lists volumes on Openstack V2 enabled hosts. The *none* agent type is the default.	
	--moving	Limit the list of volumes to volumes that are moving to a different pool.	
	[--cache_pinned {*yes*|*no*}]	Limit the list of volumes that either are (yes) or are not (no) pinned in cache.	
	[--deduped {*yes*|*no*}	Limit the list of volumes to those that have (yes) deduplicated data blocks, or volumes that do not (no) have deduplicated blocks. Volumes with deduplicated blocks include those that currently have deduplication enabled, or have had deduplication enabled in the past.	
--info	*name*	List detailed information about the specified volume.	Guest
	[--pool *pool_name*]	Display details of the volume in the specified pool. The [--verbose] option is required if the volume name is not unique within the group.	
	[--verbose]	Provide additional information about the specified volume, including metadata and the use of digests in iSCSI session.	

Option	Arguments and Suboptions	Description	Role
--create	*name*	Create a volume with the specified name.	Operator
	--size *mebibytes*	The size of the volume in mebibytes.	
	[--description *text*]	Provide a description of the volume. If there are spaces in the description, enclose the text in quotation marks.	
	[--perfpolicy *name*]	Associate a performance policy with the volume. If no performance policy is specified, the performance policy named "default" is used. The only exception to this is when a volume is created in a folder that has an --agent_type of *smis*. After the volume is created, the performance policy can be changed only to another performance policy that has the same block size as the volume.	
	[--cache_pinned {*yes*|*no*}]	If set to yes, all the contents of the volume are kept in flash cache. Pinning a volume in cache provides consistent performance guarantees for all types of workloads. The amount of flash needed to pin the volume in cache is equal to the volume quota.	
	[--thinly_provisioned {*yes*|*no*}]	Set the provisioning type of the volume to thin. The space for a new volume is provisioned immediately, and advertised to initiators that support thin provisioning. For thinly-provisioned volumes, a soft limit notification is set to initiators when the volume space usage exceeds the volume limits. The default is "yes".	
	[--limit *percent*]	Set a limit for the volume as a percentage of the volume size. If this option is not specified, the default array volume limit is used to determine the limit for this volume.	
	[--start_offline]	Start with the volume set to offline. If this option is not specified, the default is to start with the volume online.	
	[--apply_acl_to {*volume*|*snapshot*|*both*}]	Apply the specified access control list (ACL) to the volume, its snapshots, or both. If this option is not specified, the default is to apply the ACL to both the volume and its snapshots.	
	[--chapuser *user_name*]		

Option	Arguments and Suboptions	Description	Role
		Specify the name of a CHAP user who can access a volume on an iSCSI array. CHAP users can be created with the chapuser command. If this option is not specified, the default is to allow access to any CHAP user. This option is not applicable for Fibre Channel access.	
	[--initiatorgrp *group_name*]	Specify the name of the initiator group to which to restrict access. If this option is not specified, the default is no access.	
	[--lun *lun*]	Specify the logical unit number (LUN) to associate with the volume for access by Fibre Channel initiator groups. The valid range for LUNs is 0-2047. If the LUN is not specified, the system generates one. This option is not applicable for iSCSI access.	
	[--multi_initiator {*yes*|*no*}]	For iSCSI Volume Target (VST), this option specifies whether the volume and its snapshots can be accessed from multiple initiators at the same time. The default is "no". For iSCSI Group Target (GST) or Fibre Channel access, the attribute cannot be set, and vol --info always reports no.	
	[--pool *pool_name*]	Specify a storage pool to which to assign the volume. Volume data is distributed across arrays over which a specified pool is defined. If this option is not specified, the volume is assigned to the default pool.	
	[--folder *folder_name*]	Specify the name of the folder where the volume will be created. If the --pool option is specified, the folder must be present in that pool. Otherwise the folder must be present in the default pool. If the --folder is not specified, the volume is created directly in the applicable pool.	
	[--agent_type {*smis*|*none*}]	Specify the external management agent type. The default agent type is "none".	
	[--encryption_cipher {*aes-256-xts*|*none*}]		

Option	Arguments and Suboptions	Description	Role
		Specify the default cipher to use during encryption for this volume. Overriding the default value is allowed only when the --encryption_scope suboption of the group --edit command is set to "volume". If the overriding the default value is allowed, a value of "none" indicates no encryption. A value of "aes-256-xts" indicates the use of the AES-256-XTS cipher. After the volume is created and encryption is enabled, encryption on the volume cannot be changed.	
	[--dedupe_enabled {*yes*|*no*}]	Specifies whether the volume is to be deduplicated. The choice made may override the setting of the performance policy chosen for the volume.	
	[--iscsi_target_scope {*volume*|*group*}]	Specify whether the volume is to be exported under iSCSI Group Target or iSCSI Volume Target, overriding the default setting for the group.	
	[--iops_limit *iops*]	The limit for IOPS (input/output requests per second) for the Quality of Service (QoS) policy or Storage Policy-Based Management for Virtual Volumes. To set the IOPS limit to unlimited, enter "unlimited". This is the default. The minimum is 256. If both --iops_limit and --mbps_limit are specified, the --mbps_limit must not be reached before the --iops_limit.	
	[--mbps_limit *mbps*]	The limit for MiB/sec (mebibytes per second) for the Quality of Service (QoS) policy or Storage Policy-Based Management for Virtual Volumes. The lower bound is calculated as greater than or equal to the IOPS limit multiplied by the volume block size. To set the throughput limit to unlimited, enter "unlimited". The minimum throughput that can be specified is 1 MiB/sec. If this option is not specified, the throughput limit is set to the default of unlimited.	

Option	Arguments and Suboptions	Description	Role
--delete	*vol_name*	Delete an existing volume. Before you can delete any volume, you must set the volume as well as its associated snapshots offline, and disassociate any active connections. Deleting a volume that is managed by an external agent fails unless the --force option is used.	Power User
	[--pool *pool_name*]	Select the pool from which to delete the volume. The pool option is required of the volume is not unique within the group.	
	[--force]	Forcibly delete a volume that is managed by an external agent.	

Option	Arguments and Suboptions	Description	Role
--edit	*vol_name*	Modify the attributes of the volume. Specify at least one suboption.	Operator
	[--name *new_name*]	Modify the volume name. This correspondingly changes the target name associated with the volume.	
	[--description *text*]	Modify the description of the volume.	
	[--size *mebibytes*]	Modify the volume size.	
	[--perfpolicy *name*]	Modify the performance policy that is associated with the volume. The new performance policy must have a block size that matches the block size of the performance policy already associated with the volume.	
	[--cache_pinned {*yes*|*no*}]	Modify whether to keep all the contents of this volume in flash cache. This allows for consistent performance guarantees for all types of workloads. The amount of flash needed to pin the volume is equal to the quota for the volume.	
	[--readonly {*yes*|*no*}]	Enter *yes* to enable read-only access to the volume, or *no* to enable read and write access to the volume.	
	[--force]	Forcibly set the volume to read-only when there are connected initiators, or when reducing the size of the volume.	
	[--thinly_provisioned {*yes*|*no*}]	Modify the provisioning type of the volume. The space for the volume is provisioned immediately, but for advertising status, the change takes effect only for new connections to the volume. Existing initiators must disconnect and reconnect for the new setting to take effect. A soft limit notification is set to initiators when the volume space usage exceeds the volume limits.	
	[--limit *percent*]	Modify the volume limit.	
	[--multi_initiator {*yes*|*no*}]	For the iSCSI access protocol, this option specifies whether the volume and its snapshots can be accessed from multiple initiators at the same time. For the Fibre Channel access protocol, this attribute cannot be set, and vol --info always reports "yes".	
	[--agent_type {*smis*|*none*}]	Modify the external management agent type.	
	[--dedupe_enabled {*yes*|*no*}]		

Option	Arguments and Suboptions	Description	Role
		Modify whether the volume is to be deduplicated. The setting overrides the settings of the performance policy associated with the volume.	
	[--iscsi_target_scope {*volume*|*group*}]	Modify whether the volume is to be exported under iSCSI Group Target or iSCSI Volume Target.	
	[--iops_limit *iops*]	The limit for IOPS (input/output requests per second) for the Quality of Service (QoS) policy or Storage Policy-Based Management for Virtual Volumes. To set the IOPS limit to unlimited, enter "unlimited". This is the default. The minimum is 256 requests. If both --iops_limit and --mbps_limit are specified, the --mbps_limit must not be reached before the --iops_limit.	
	[--mbps_limit *mbps*]	The limit for MiB/sec (mebibytes per second) for the Quality of Service (QoS) policy or Storage Policy-Based Management for Virtual Volumes. The lower bound is calculated as greater than or equal to the IOPS limit multiplied by the volume block size. To set the limit to unlimited, enter "unlimited". The minimum throughput that can be specified is 1 MiB/sec. If this option is not specified, then the throughput limit is set to the default of unlimited.	
	[--pool *pool_name*]	Edit the volume in the specified pool. The --pool option is required if the volume name is not unique within the group.	
--online	*volume_name*	Set the volume online. Setting a volume online allows initiators to connect to it and allows writes.	Operator
	[--pool *pool_name*]	Online the volume in the specified pool. The --pool option is required if the volume is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--offline	*volume_name*	Set the volume offline. Setting a volume offline prevents initiators from connecting to it. Setting the volume offline fails if there are open connections from initiators or a virtual volume is bound to a protocol endpoint, unless the --force is used.	Operator
	--force	Forcibly set the specified volume offline, even if there are open connections from initiators. This action effectively disables the volume.	
	[--pool *pool_name*]	Offline the volume in the specified pool. The --pool option is required if the volume is not unique within the group.	
--snap	*vol_name*	Snapshot the specified volume.	Operator
	--snapname *name*	Specify the name of the snapshot to create.	
	[--description *text*]	Provide a description of the snapshot. If there are spaces in the description, enclose the text in quotation marks.	
	[--start_online]	Start the snapshot online. By default, snapshots are set offline when they are created.	
	[--allow_writes]	Allow applications to write to the snapshot. By default, writes are disallowed.	
	[--pool *pool_name*]	Pool in which the volume resides. The --pool option is required if the volume is not unique within the group.	
--restore	*vol_name*	Restore a volume from one of its snapshots. The volume must be taken offline before it can be restored. A snapshot of the volume is taken before restoring the volume to facilitate restoring to an earlier point in time, if needed.	Power User
	--snapname *snap_name*	Specify the snapshot from which to restore the volume.	
	[--pool *pool_name*]	Select thee volume in the specified pool to be restored. the --pool option is required if the volume name is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--clone	*vol_name*	Clone a volume from one of its snapshots to create a new volume. You can use clones to restore corrupt volumes or to create an environment for testing changes.	Operator
	--snapname *snap_name*	Name of the snapshot to use to create the clone.	
	--clonename *clone_name*	The name for the new, cloned volume.	
	[--folder *folder_name*]	The name of the folder in which the cloned volume will reside. The folder must be in the same pool as the original volume. If this option is not specified, the cloned volume is placed in the same folder as the original volume.	
	[--description *text*]	Provide a description of the cloned volume. If there are spaces in the description, enclose the text in quotation marks.	
	[--readonly {*yes*|*no*}]	Set the clone to be a read-only or read/write volume.	
	[--thinly_provisioned {*yes*|*no*}]	Set the provisioning type of the clone to thin. The space for the clone is provisioned immediately and advertised to initiators that support thin provisioning. The default is "yes".	
	[--limit *percent*]	Specify the limit for the cloned volume as a percentage of the volume size.	
	[--start_offline]	Set the cloned volume offline. By default, clones are set online when you create them.	
	[--apply_acl_to {*volume*|*snapshot*|*both*}]	Apply a specified access control list to the cloned volume, its snapshots, or both. If this option is not specified, the default is to apply the ACL to both.	
	[--chapuser *user_name*]	Specify the name of a CHAP user who can access this cloned volume on an iSCSI array. If this option is not specified, the default is to allow access to any CHAP user. This option does not apply to volumes with Fibre Channel access.	
	[--initiatorgrp *group_name*]	Specify the name of the initiator group to which to restrict access. If this option is not specified, access is allowed to any initiator group.	
	[--lun *lun*]		

Option	Arguments and Suboptions	Description	Role
		Specify the logical unit number (LUN) to associate with the cloned volume for access by Fibre Channel initiator groups. The valid range of values for LUN is 0-2047. If this option is not specified, the system generates one. This option is not applicable to volumes with iSCSI access.	
	[--multi_initiator {*yes*|*no*}]	For iSCSI access, this option specifies whether the volume clone and its snapshots can be accessed from multiple initiators at the same time. The default is to match the setting to that of the parent volume. For Fibre Channel access, this attribute cannot be set, and vol --info always reports "no".	
	[--cache_pinned {*yes*|*no*}]	If set to *yes*, all contents of the cloned volume are kept in flash cache. This allows for consistent performance guarantees for all types of workloads. The amount of flash needed to pin the cloned volume is equal to the limit for the volume plus the size of the snapshot it is cloned from.	
	[--dedupe_enabled {*yes*|*no*}]	If set to *yes*, the cloned volume uses the deduplication setting specified, instead of inheriting the deduplication setting from its parent volume.	
	[--iops_limit *iops*]	The limit for IOPS for the QoS policy or Storage Policy-Based Management for Virtual Volumes for the cloned volume. The default value is the IOPS limit of the parent volume. To set the IOPS limit to unlimited, enter "unlimited". The minimum is 256. If both --iops_limit and --mbps_limit are specified, the --mbps_limit must not be reached before the --iops_limit.	
	[--mbps_limit *mbps*]	The throughput limit for the cloned volume in mebibytes (MiB/sec). To set the throughput limit to unlimited, use "unlimited". The default is the throughput limit of the parent volume. The minimum throughput value is 1.	
	[--pool *pool_name*]	Select the volume to be cloned from the specified pool. The --pool option is required if the volume name is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--assoc	*vol_name*	Associate a volume with a volume collection.	Operator
	--volcoll *name*	The volume collection with which to associate the volume.	
	[--pool *pool_name*]	Associate a volume in the specified pool with the volume collection. The --pool option is required if the volume name is not unique within the group.	
--dissoc	*vol_name*	Disassociate a volume from a volume collection. When a volume is disassociated from a volume collection, any snapshots of the volume are not deleted when the corresponding snapshot collection is deleted, either manually or per-retention-policy for the protection schedule by which they are managed.	Operator
	[--pool *pool_name*]	Disassociate a volume in the specified pool from a volume collection. The --pool option is required if the volume name is not unique within the group.	
	--force	Forcibly disassociate a volume managed in the context of a vCenter volume collection.	
--addacl	*volume_name*	Add an access control list (ACL) record to the volume.	Operator
	--apply_acl_to {*volume*|*snapshot*|*both*}	Apply the access control list to the volume, its snapshots, or both. If this option is not specified, the default action is to apply the ACL to both the volume and its snapshots.	
	[--chapuser *user_name*]	Specify the CHAP user to restrict access to. This option does not apply to volumes with Fibre Channel access.	
	[--initiatorgrp *group_name*]	Name of the initiator group to which to restrict access.	
	[--lun *lun*]	Specify the logical unit number (LUN) to associate with this volume for access by a Fibre Channel initiator group. Valid LUNs are in the 0-2047 range. If not specified, the system generates a number. This option does not apply to volumes with iSCSI access.	
	[--pool *pool_name*]	Add an access control record to the volume in the specified pool. The --pool option is required if the volume name is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--removeacl	*vol_name*	Remove an access control list record from the specified volume.	Operator
	--apply_acl_to {*volume*|*snapshot*|*both*}	Remove the access control list record from the volume, its snapshots, or both. If this option is not specified, the default is to remove it from both.	
	[--chapuser *user_name*]	Remove the access control record for the specified CHAP user. This option does not apply to volumes with Fibre Channel access.	
	[--initiatorgrp *group_name*]	Remove the access control record for the specified initiator group.	
	[--pool *pool_name*]	Remove access control record from the volume in the specified pool. The --pool option is required if the volume name is not unique within the group.	
--claim	*vol_name*	Claim ownership of the specified volume. Claiming ownership of a volume allows the new owner to modify volume attributes. For example, when a replica volume is disassociated on an upstream array, the volume ownership for the replica stays with the upstream array. To modify the volume on the local system, claim ownership. Note that **--claim** works on volumes that are not in a volume collection. For volumes in a collection, use **volcoll --promote**. **Note:** The claim command does not apply to synchronous replication, because both the source and destination volumes are in the same group.	Power User
	[--pool *pool_name*]	Claim the volume in the specified pool. The --pool option is required if the volume name is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--move	*vollist*	Move the specified list of volumes, any related volumes and their snapshots, to the specified destination pool. Related volumes are those that share data blocks, and therefore include any volume which has a clone or parent relationship with the specified volumes, either directly or indirectly. Multiple volumes may be specified using a comma-delimited list. Volume moves can take a long time, depending on the number of volumes and snapshots being moved, and the overall data use across the volumes and snapshots. Use vol --info to monitor the move. Use the --assign_folder option to assign a volume to a folder in the same pool.	Power User
	--dest_pool *pool_name*	The name of the destination pool for the volumes. The destination pool must have sufficient space to accommodate the current usage and reserves of all the volumes being moved. The destination pool also must be under the limit for the number of volumes in a pool.	
	[--dest_folder *folder_name*]	The name of the destination folder for the volumes. If not specified, the volumes will be contained directly by the destination pool.	
	[--pool *pool_list*]	List of source pools corresponding to the volumes in the volume list. This option is required if one or more volumes in the volume list are not unique within the group. If the pool list is specified, the count of entries in the pool list must equal the count of entries in the volume list. Multiple pools can be specified using a comma-separated list.	
--abort_move	*vol_name*	Stop an in-progress move of the specified volume to another pool. Any data which has already been moved to the destination pool would need to be moved back to the source pool. Once this operation is initiated, it can take a long time, depending on the amount of data transferred to the destination pool before the move was stopped.	Power User
	[--pool *pool_name*]	Source pool of the volume that was in the process of moving. The --pool option is required if the volume name is not unique within the group.	

Option	Arguments and Suboptions	Description	Role
--assign_folder	*vol_name*	Assign a volume to a folder in the same pool. If the volume is already in the specified folder, no action is taken. If no folder is specified, the folder assignment of the volume is removed. Use the --move option to move a volume and any related volumes to another pool.	
	--folder *folder_name*	Name of the folder. The folder must be in the same pool as the volume.	
	[--pool *pool_name*]	Pool in which the volume resides. The --pool option is required if the volume name is not unique within the group.	

Examples

This example lists all volumes on an array.

```
Nimble OS $ vol --list
--------------------+----------+------+----------+----------+---------+-------+---------
Name                  Size       Online Offline     Usage      Reserve % Limit
 % Path
                      (MB)               Reason      (MiB)
--------------------+----------+------+----------+----------+---------+-------+---------
AppAD01                87041 No     Replica      11179       0.00     100.00
 default:/
AppE2K7Node2          153600 No     Replica      63550       0.00     100.00
 default:/
EX03-Exchange         512000 Yes    N/A          63325       0.00     100.00
 default:/
EX03-ExchangeLog      102400 Yes    N/A          23049       0.00     100.00
 default:/
itbackup             1048576 Yes    N/A         301628       0.00     100.00
 default:/
share                2097152 Yes    N/A         781912       0.00     100.00
 default:/
vcenterbackup         102400 Yes    N/A          41024       0.00     100.00
 default:/
```

This example shows detailed information about a volume named *public* that is configured on an iSCSI array. The output is truncated to save space.

```
Nimble OS $ vol --info public
Name: public
Serial number: 38c46c34d60faf846c9ce9058d6e41a2
iSCSI target: iqn.2007-11.com.nimblestorage:public-v354fae.0000f.a241
Description: public volume
Owned by: greyful
Size (MB): 2097152
Pool: default
Move to pool: N/A
Move aborting: N/A
Move data migrated: N/A
Move data remaining: N/A
Move start time: N/A
Move estimated completion time: N/A
Performance policy: Windows File Server
```

```
Block size (bytes): 4096
Reserve: 0.00%
Warn level: 80.00%
Quota: 100.00%
Snapshot reserve: 0.00%
Snapshot warn level: N/A
Snapshot quota: unlimited
Snapshot count: 2
...
...
Agent Type: none
Online Snapshots:
Encryption cipher: aes-256-xts
```

This example shows detailed information about a volume named *semiprivate* that is configured on a Fibre Channel array. The output is truncated to save space.

```
Nimble OS $ vol --info semiprivate
Name: semiprivate
Serial number: 4e136d3f091130236c9ce9006547755f
Target name: 56:c9:ce:90:b9:f2:ca:00
Description:
Owned by: corp
Size (MB): 16777216
Pool: default
Move to pool: N/A
Move aborting: N/A
Move data migrated: N/A
Move data remaining: N/A
Move start time: N/A
Move estimated completion time: N/A
Performance policy: default
Block size (bytes): 4096
Reserve: 0.00%
Warn level: 80.00%
Quota: 100.00%
Snapshot reserve: 0.00%
Snapshot warn level: N/A
Snapshot quota: unlimited
Snapshot count: 0
Volume usage (MB): 35331
Volume compression: 2.65X
Volume space saved (MB): 58204
...
...
        Access Protocol: fc
        LUN: 0
Connected Initiators:
 Initiator: A11_hba3 (21:00:00:0e:1e:19:60:91)
       Target FC Interface: (array7 : B : fc6.1) (56:c9:ce:90:4f:a6:c8:06)
       ALUA: standby
       PR Key: 0
       Target FC Interface: (array7 : B : fc10.1) (56:c9:ce:90:4f:a6:c8:08)
       ALUA: standby
       PR Key: 0
       Target FC Interface: (array7 : A : fc6.1) (56:c9:ce:90:4f:a6:c8:02)
       ALUA: active/optimized
       PR Key: 0
       Target FC Interface: (array7 : A : fc10.1) (56:c9:ce:90:4f:a6:c8:04)
       ALUA: active/optimized
       PR Key: 0
```

```
...
...
```

This example creates a volume named *publicvol* with a quota of 80 percent and no reserve. This is in case the volume must be expanded later. No ACLs are included. They are added later.

```
Nimble OS $ vol --create  publicvol --size 30
--description "for use by everyone" --quota 80 --warn_level 75
```

This example modifies the size and description of a volume named *publicvol* and sets the volume to read only.

```
Nimble OS $ vol --edit publicvol --size 50 --description "for use by me"
--readonly yes
```

This example deletes the volume named *publicvol*.

```
Nimble OS $ vol --delete publicvol
```

This example brings a volume named *publicvol* online.

```
Nimble OS $ vol --online publicvol
```

This example takes the volume named *publicvol* offline.

```
Nimble OS $ vol --offline publicvol --force
```

This example creates a snapshot of the volume named *publicvol* and makes the new *publicvolsnap* snapshot writable.

```
Nimble OS $ vol --snap publicvol
--snapname publicvolsnap --description "a snapshot of the public volume"
--allow_writes
```

This example takes a volume named *publicvol* offline, restores it from the snapshot named *publicvolsnap*, and brings the restored volume online.

```
Nimble OS $ vol --offline publicvol

Nimble OS $ vol--restore publicvol --snapname publicvolsnap

Nimble OS $ vol --online publicvol
```

This example lists snapshots, clones a volume named *publicvol* from its snapshot named *publicvolsnap*, and then verifies the success of the operation.

```
Nimble OS $ vol --list
--------------------+----------+------+-----------+----------+---------+-------
Name                  Size       Online Offline     Usage      Reserve % Quota
 %
                      (MB)               Reason      (MB)
--------------------+----------+------+-----------+----------+---------+-------
publicvol                    30 Yes    N/A                 0       0.00   80.00

Nimble OS $ snap --list
--------------------+--------------------+----------+------+-------+-----------
Volume                Snapshot             Size       Online Status  New Data
Name                  Name                 (MB)                      (MB)
--------------------+--------------------+----------+------+-------+-----------
```

```
publicvol            publicvolsnap                30 No     Okay
 0

Nimble OS $ vol --clone publicvol --snapname publicvolsnap
--clonename publicclone --description "testing version compatibility"

Nimble OS $ vol --list
--------------------+----------+------+-----------+----------+---------+-------
Name                Size       Online Offline     Usage      Reserve % Quota
 %
                     (MB)              Reason      (MB)
--------------------+----------+------+-----------+----------+---------+-------
publicclone                  30 Yes    N/A                0      0.00   100.00
publicvol                    30 Yes    N/A                0      0.00    80.00
```

This example lists available volume collections and then associates one of the collections with a volume named *publicvol*.

```
Nimble OS $ volcoll --list
--------------------+---------------+----------------------------------------
Volume Collection   Application     Owned By
Name                Synchronization
--------------------+---------------+----------------------------------------
volcollone          none            greyhound

Nimble OS $ vol --assoc publicvol --volcoll volcollone
```

This example disassociates a volume named *publicvol* from a volume collection.

```
Nimble OS $ vol --dissoc publicvol
```

This example lists available CHAP users on an iSCSI array, adds an ACL for the CHAP user to a volume named *publicvol*, and then shows the ACL definitions. The vol --info output is truncated to save space and prevent line wrapping.

```
Nimble OS $ chapuser --list
-----------------------------+--------------------------------------------------
CHAP User Name                 Password
-----------------------------+--------------------------------------------------
nimblechap                     nimblechapuser

Nimble OS $ vol --addacl publicvol --apply_acl_to volume
--chapuser nimblechap

Nimble OS $ vol --info publicvol
Name: publicvol
Serial number: ff3161997a340eb76c9ce9004ec1399f
iSCSI target: iqn.2007-11.com.nimblestorage:publicvol-v5e0c4565c135842b...
Description: for use by everyone
...
...
Access Control List:
        Apply to: volume & snapshot
        Initiator Group: *
        CHAP user: *
        Apply to: volume
        Initiator Group: *
        CHAP user: nimblechap
Connected Initiators:
...
...
```

This example removes the name of a CHAP user from the access control record on a volume named *publicvol* on an iSCSI array.

```
Nimble OS $ vol --removeacl publicvol --apply_acl_to volume --chapuser nimblechap
```

This example claims ownership of a volume named *publicvol* for a replication partner named greyhound.

```
Nimble OS $vol --claim publicvol --partner greyhound
```

This example moves a volume named *publicvol* to a storage pool named *secondary*.

```
Nimble OS $vol --move  publicvol --dest_pool secondary
```

Use the **--moving** option to view the status of volumes in the process of moving.

This example stops the in-process move of a volume named *publicvol*.

```
Nimble OS $vol --abort_move  publicvol
```

volcoll

The volcoll command and its options manage volume collections. Volume collections are logical groups of volumes that share protection characteristics, such as snapshot and replication schedules. Volume collection names are case-sensitive.

You can create volume collections or base them on predefined protection templates. Snapshots for all volumes in a collection are captured synchronously to ensure that the data across these volumes is mutually consistent. Volumes in a collection share snapshot and replication schedules as well as the retention policies for those snapshots. Similarly, during disaster recovery, all volumes in a volume collection simultaneously fail over to the replica array as a group. You need to create a volume collection for each application.

Volume collections may include multiple schedules. When schedules overlap, snapshots, or replicas, are created for every schedule. Schedules affect all volumes that are assigned to the volume collection. Volume collections can have up to two downstream replication partners.

Synopsis

volcoll --help

volcoll --list
[--replication_type {*periodic_snapshot* | *synchronous*}]

volcoll --info *name*

volcoll --create *name*
--prottmpl *name*
[--description *text*]
--app_sync {*none* | *vss* | *vmware* | *generic*} [--app_server *server*]
[--app_id {*exchange* | *sql2005* | *sql2008* | *sql2012* | *exchange_dag* | *sql2014* | *sql2016* | *hyperv*}]
[--app_cluster_name *cluster_name*]
[--app_service_name *service_name*]
[--vcenter_hostname *server*]
[--vcenter_username *user_name*]
[--vcenter_password *password*]
[--replication_type {*periodic_snapshot* | *synchronous*}]
[--agent_hostname *server*]
[--agent_username *user_name*]

[--agent_password *password*]

volcoll --delete *name*

volcoll --edit *volcoll_name*
[--name *new_name*]
[--description *text*]
[--app_sync {*none* | *vss* | *vmware* | *generic*}]
[--app_server *server*]
[--app_id {*exchange* | *sql2005* | *sql2008* | *sql2012* | *exchange_dag* | *sql2014* | *sql2016* | *hyperv*}]
[--app_cluster_name *cluster_name*]
[--app_service_name *service_name*]
[--vcenter_hostname *server*]
[--vcenter_username *user_name*]
[--vcenter_password *password*]
[--agent_hostname *server*]
[--agent_username *user_name*]
[--agent_password *password*]

volcoll --delete *name*

volcoll --validate *name*

volcoll --addsched *name*
--schedule *name*
--repeat *period*
--repeat_unit {*minutes* | *hours* | *days* | *weeks*}
[--at *time*]
[--until *time*]
[--days {*all|day1,day2,...*}]
[--retain {*number*}]
[--replicate_to *partner*]
[--replicate_every *number*]
[--num_retain_replica *number*]
[--alert_threshold *hh:mm*]
[--snap_verify {*yes* | *no*}]
[--skip_db_consistency_check {*yes* | *no*}]
[--disable_appsync {*yes* | *no*}]
[--external_trigger {*yes* | *no*}]

volcoll --editsched *volcoll_name*
--schedule *name*
[--name *new_name*]
[--repeat *period*]
[--repeat_unit {*minutes* | *hours* | *days* | *weeks*}]
[--at *time*]
[--until *time*]
[{*all|day1,day2,...*}]
[--retain *number*]
[--replicate_to *partner*]
[--replicate_every *number*]
[--num_retain_replica *number*]
[--use_downstream_for_DR]
[--alert_threshold *hh:mm*]
[--snap_verify {*yes* | *no*}]
[--skip_db_consistency_check {*yes* | *no*}]

[--disable_appsync {*yes* | *no*}]
[--external_trigger {*yes* | *no*}]
[--non-interactive]
[--force]

volcoll --deletesched *name*
--schedule *name*
[--force]

volcoll --snap *name*
--snapcoll_name *name*
[--description *text*]
[--start_online]
[--allow_writes]
[--replicate]
[--replicate_to *partner*]]
[--disable_appsync {*yes* | *no*}]
[--external_trigger {*yes* | *no*}]
[--snap_verify {*yes* | *no*}]
[--skip_db_consistency_check {*yes* | *no*}]

volcoll --promote *name*

volcoll --demote *volcoll*
--partner *name*
[--non_interactive]

volcoll --handover *name*
--partner *name*
[--no_reverse]
[--abort]
[--override_upstream_down]
[--non_interactive]

volcoll--stop_repl *name*
--schedule *name*

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Guest
--list	[--replication_type {*periodic_snapshot* | *synchronous*}]	List volume collections.	Guest
--info	*name*	Provide detailed information about the specified volume collection, type of replication, and its schedules. Volume collection names are case-sensitive.	Guest

Option	Arguments and Suboptions	Description	Role
--create	*name*	Create a volume collection.	Operator
	--prottmpl *name*	Use a specific protection template as a basis for defining protection-related attributes of the volume collection. If other options are specified, those options then override attributes that have been derived from a protection template. Any changes to the protection template that occur after the volume collection is created do not affect any protection attributes of this collection.	
	[--description *text*]	Enter a plain-language description of the volume collection. If there are spaces in the description, enclose the text in quotation marks.	
	--app_sync {*none|vss|vmware|generic*}	Define the application synchronization to apply. If you specify VSS synchronization, supply the --app_* credentials. If you specify VMware synchronization, supply the --vcenter_* credentials. If you have created a custom agent, specify the generic app type when creating a volume collection.	
	[--app_server *server*]	Specify the application server when VSS synchronization is enabled.	
	[--app_id {*exchange|sql2005|sql2008|sql2012|exchange_dag|sql2014|sql2016|hyperv*}]	Define the application running on the server when VSS synchronization is enabled, including MS-Hyper-V. Exchange DAG specifies that it is a Microsoft Exchange Server 2010 or later supporting Database Available Group.	
	[--app_cluster_name *cluster_name*]	Specify the cluster name if the application that uses VSS synchronization is running within a Windows clustering environment.	
	[--app_service_name *service_name*]	Specify the instance name of the	
	[--vcenter_hostname *server*]	Specify the host name of the vCenter server with which the Nimble array will communicate when VMware synchronization is enabled. A custom port number can be specified with vCenter using "**:**".	
	[--vcenter_username *user_name*]	Specify the Windows user name to use to create a Nimble account on the vCenter server when VMware synchronization is enabled.	
	[--vcenter_password *password*]	Specify the password associated with the user name when VMware synchronization is enabled.	

Option	Arguments and Suboptions	Description	Role
	[--replication_type {*periodic_snapshot* | *synchronous*}]	Specify the replication type. The default is periodic snapshot.	
	--agent_hostname *server*	Specify a backup agent host name. A customer port number can be specified with vCenter using "**:**".	
	--agent_username *user_name*	Specify a backup agent user name.	
	--agent_password *password*	Specify the password associated with the backup agent user name.	

Option	Arguments and Suboptions	Description	Role
--edit	*volcoll_name*	Modify the specified volume collection. New volume collections based on the template are affected by the modifications, but existing volume collections are not. Specify at least one suboption. **Note:** For synchronous replication, you cannot convert a volume collection that was configured from one form of replication to another form of replication.	Operator
	[--name *new_name*]	Modify the name of the volume collection.	
	[--description *text*]	Modify the plain-language description of the volume collection.	
	--repl_priority {*normal|high*}	Modify the replication priority level.	
	--app_sync {*none|vss|vmware*}	Modify the application synchronization to apply.	
	[--app_server *server*]	Modify the application server for VSS synchronization.	
	[--app_id {*exchange|sql2005|sql2008|sql2012|exchange_dag|sql2014|sql2016|hyperv*}]	Modify the application running on the server for VSS synchronization, including MS-Hyper-V.	
	[--app_cluster_name *cluster_name*]	Modify the cluster name if the application that uses VSS synchronization is running within a Windows clustering environment.	
	[--app_service_name *service_name*]	Modify the instance name of the service if the application that uses VSS synchronization is running within a Windows clustering environment.	
	[--vcenter_hostname *server*]	Modify the host name of the vCenter server with which the Nimble array communicates for VMware synchronization.	
	[--vcenter_username *user_name*]	Modify the Windows user name for VMware synchronization.	
	[--vcenter_password *password*]	Modify the password associated with the user name for VMware synchronization.	
	[--agent_hostname *server*]	Modify the backup agent host name.	
	[--agent_username *user_name*]	Modify the backup agent user name.	
	[--agent_password]	Modify the password associated with the backup agent user.	

Option	Arguments and Suboptions	Description	Role
--delete	*name*	Delete the specified volume collection. Disassociate all volumes from the volume collection using the vol --dissoc before deleting the volume collection.	Power User
--validate	*name*	Validate a volume collection with either Microsoft VSS, VMware, or generic application synchronization.	

Option	Arguments and Suboptions	Description	Role
--addsched	*name*	Create a protection schedule and add it to the specified volume collection.	Operator
	--schedule *name*	Specify the name of the snapshot schedule to associate with this volume collection.	
	--repeat *period*	Specify the frequency of snapshots. If you do not specify this option, the default of *1* is used.	
	--repeat_unit {*minutes|hours|days|weeks*}	Specify the unit of time for --repeat to define the frequency of snapshots. If you do not specify this option, the default unit of *days* is used.	
	[--at *time*]	Specify the time of day to start taking snapshots. If you do not specify this option, the default value of *12:00* (midnight) is used.	
	[--until *time*]	Specify the time of day to stop taking snapshots. This option is applicable only when the repeat frequency is more than one snapshot per day. If you do not specify this option, the default value of *11:59* (for snapshots taken every minute) is used.	
	[--days {*all|day1,day2,...*}]	Specify the days on which to take snapshots unless the repeat unit is *weeks*. You can enter a comma-separated list of days, for example, Mon,Wed,Fri. If you do not specify this option, the default value of *all* is used.	
	[--retain *number*]	Specify the maximum number of snapshots to keep. If the schedule includes replication, the array retains the latest replicated snapshot, which may exceed the specified retention value. This is necessary to ensure efficient replication performance.	
	[--replicate_to *pool_partner*]	Specify the replication partner for replicated snapshots. Use the partner --create command on both arrays to configure replication partners.	
	[--replicate_every *number*]	Specify which snapshots assigned to the schedule should be replicated. For example, setting this to 5 replicates every fifth snapshot. If snapshots are replicated and you do not specify this option, all snapshots are replicated.	
	[--num_retain_replica *number*]	Specify the number of snapshots to retain on the replication partner.	

Option	Arguments and Suboptions	Description	Role
	[--alert_threshold *hh:mm*]	If replicating a snapshot takes more than this amount of time to complete, an alert will be generated. If you do not specify this option, the default of 24 hours is used. Enter *00:00* to disable this alert.	
	[--snap_verify {*yes*|*no*}]	Specify whether to run a verification tool on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization. The tool used to verify snapshots depends on the type of application. For example, if the application synchronization is VSS and the application ID is *exchange*, the eseutil tool is run on the snapshots. If verification fails, the logs are not truncated.	
	[--skip_db_consistency_check {*yes*|*no*}]	Specify whether to skip consistency checks for database files on snapshots created by this schedule. Use this option with snapshot schedules of a volume collection that has application synchronization of VSS, the application ID is *exchange_dag*, --snap_verify is set to *yes*, and --disable_appsync is set to *no*. Skipping consistency checks is only recommended if each database in a DAG has multiple copies.	
	[--disable_appsync *{yes|no}*]	Specify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	
	[--external_trigger *{yes|no}*]	Specify whether to use an externally driven schedule with no internal timers to create manual snapshots.	

Option	Arguments and Suboptions	Description	Role
--editsched	*volcoll_name*	Modify the specified schedule in the specified volume collection. Specify at least one suboption in addition to --schedule *name*.	Operator
	--schedule *name*	Specify the name of the snapshot schedule to modify.	
	[--name *new_name*]	Modify the name of the snapshot schedule.	
	[--repeat *period*]	Modify the frequency of snapshots.	
	[--repeat_unit {*minutes|hours|days|weeks*}]	Modify the unit of time for --repeat to define the frequency of snapshots.	
	[--at *time*]	Modify the time of day to start taking snapshots.	
	[--until *time*]	Modify the time of day to stop taking snapshots.	
	[--days {*all|day1,day2,...*}]	Modify the days on which to take snapshots unless the repeat unit is *weeks*.	
	[--retain *number*]	Modify the maximum number of snapshots to keep.	
	[--replicate_to *partner*]	Modify the replication partner for replicated snapshots.	
	[--replicate_every *number*]	Modify the number of snapshots assigned to the schedule should be replicated.	
	[--num_retain_replica *number*]	Modify the number of snapshots to retain on the replication partner.	
	[use_downstream_for_DR]	Break synchronous replication for the specified volume collection and present downstream volumes to the host or hosts. Use only if necessary during disaster recovery.	
	[--alert_threshold *hh:mm*]	Modify the replication alert threshold. If a snapshot takes more than this amount of time to complete, and alert is generated. Enter *00:00* to disable the alert.	
	[--snap_verify {*yes|no*}]	Modify whether to run a verification tool on snapshots created by this schedule.	
	[--skip_db_consistency_check *{yes|no}*]	Modify whether to skip consistency checks for database files on snapshots created by this schedule.	
	[--disable_appsync {*yes|no*}]	Modify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	
	[--external_trigger {*yes|no*}]		

Option	Arguments and Suboptions	Description	Role
		Modify whether to use an externally driven schedule with no internal timers to create manual snapshots.	
	[--non_interactive]	Do not prompt for confirmation.	
	[--force]	Forcibly edit a protection schedule in a volume collection managed by vCenter.	
--deletesched	*name*	Delete a protection schedule of specified volume collection. Existing snapshots created by this schedule are unaffected and will no longer be subject to retention or replication scheduling.	Operator
	--schedule *name*	Specify the snapshot schedule to delete.	
	[--force]	Forcibly delete a protection schedule of a volume collection managed by vCenter.	

Option	Arguments and Suboptions	Description	Role
--snap	*name*	Create a snapshot of volumes that are associated with the specified volume collection. All created snapshots are consistent with each other. If the volume collection is application synchronized (VMware or VSS), the volume snapshots are synchronized with the application as well. If application synchronization is disabled for all protection schedules in the collection, the snapshot is not taken and an error occurs.	Operator
	--snapcoll_name *name*	Specify the name of the snapshot collection that is being created. The *name* is used for the snapshot as well as the snapshot collection.	
	[--description *text*]	Supply a plain-language description of the snapshot collection. If there are spaces in the description, enclose the text in quotation marks.	
	[--start_online]	Specify whether to start the snapshot collection in an online state. If you do not specify this option, the snapshot collection is started offline.	
	[--allow_writes]	Specify whether to allow applications to write to the new snapshot collection. If you do not specify this option, the snapshot collection disallows writes.	
	[--replicate]	Specify whether to replicate the snapshot collection and its members to a replication partner. If you do not specify this option, the snapshot collection is not replicated. If you do specify this option, create the replication partner using the partner --create command.	
	[--replicate_to *partner*]	Enter the name of the partner to which you want to replicate the newly created snapshot collection. This must be one of the partners associated with the volume collection. If the schedule for the volume collection has multiple partners, you must include this option.	
	[--disable_appsync {*yes*|*no*}]	Specify whether to disable application-synchronized snapshots and create crash-consistent snapshots instead.	
	[--external_trigger {*yes*|*no*}]	Specify whether to use an externally driven schedule to take the snapshot instead of the Nimble-driven schedule.	
	[--snap_verify {*yes*|*no*}]	Specify whether to run a verification tool on this snapshot collection.	

Option	Arguments and Suboptions	Description	Role
	[--skip_db_consistency_check {*yes*|*no*}]	Specify whether to skip consistency checks for database files in this snapshot collection.	
--promote	*name*	Take ownership of the specified volume collection. This option is used on a downstream replication partner. The volumes that are associated with the volume collection are set to online, so they are available for reading and writing. Replication is disabled on the affected schedules and must be reconfigured if needed. Snapshot retention for the affected schedules is set to the greater of the current local or replica retention values. **Note:** You cannot use this command with synchronous replication.	Power User
--demote	*name*	Release ownership of the specified volume collection to the specified replication partner. The volumes associated with the volume collection are set to offline and a snapshot is created. Full control over the volume collection is then transferred to the new owner. You can use this option following --promote to revert the volume collection to its previously configured state. This operation does not modify the configuration on the new owner array. The new owner array must be running to obtain its identity information. **Note:** You cannot use this command with synchronous replication.	Power User
	--partner *name*	Specify the replication partner that is given ownership of the volume collection.	
	[--non-interactive]	Do not prompt for confirmation of this action.	

Option	Arguments and Suboptions	Description	Role
--handover	*name*	Gracefully transfer ownership of the specified volume collection to the specified replication partner. Ownership and full control is given to the replication partner. The volumes that are associated with the volume collection are taken offline before the final snapshot is taken and replicated. Therefore, full data synchronization is ensured as part of the transfer.	Power User
	--partner *name*	Enter the name of the downstream partner. Ownership will be transferred to this replication partner.	
	[--no_reverse]	Do not automatically reverse the direction of replication. If you do not specify this option, the new owner begins to replicate to the original array when the handover completes.	
	[--abort]	Cancel an in-progress handover.	
	[--override_upstream_down]	Allows a handover even if the upstream site is down.	
	[--non-interactive]	Do not prompt for confirmation of this action.	
--stop_repl	*name*	Discontinue replication for a protection schedule of the specified volume collection.	Power User
	--schedule *name*	Specify the protection schedule to discontinue.	

Examples

This example lists the volume collections on the array.

```
Nimble Storage $ volcoll --list
--------------------+---------------+--------------------
Volume Collection    Application     Owned By
Name                 Synchronization
--------------------+---------------+--------------------
volcoll-DB1           VSS             array1
daily                 none            array7
```

This example lists the periodic snapshot volcolls.

```
Nimble Storage $ volcoll --list --replication_type periodic_snapshot
--------------------+---------------+-------------------------------------------
Volume Collection    Application     Owned By
Name                 Synchronization
--------------------+---------------+-------------------------------------------
periodic-vc1         none            g1a1
periodic-vc2         none            g1a1
```

```
periodic-vc3         none            g1a1
```

This example shows detailed information about a volume collection named *daily*.

```
 # volcoll --info daily
Name: daily
Description:
Owned by: array7
Application synchronization: none
Application server: N/A
Application ID: N/A
Cluster name: N/A
Service name: N/A
VMware vCenter hostname: N/A
VMware vCenter username: N/A
VMware vCenter password: N/A
Associated volumes: na-Vol1, na-Vol5, na-Vol4, na-Vol3
Associated pinned volumes: none
Snapshot collection count: 1
Created: Feb 19 2014 15:30:50
Last configuration change: Feb 19 2014 17:06:25
Schedule Name: daily
        Description:
        Type: Nimble Schedule
        Owned by: array7
        Repeats: 1 day(s)
        At: 10:38:00 a.m.
        Until: N/A
        Days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday
        Last snapshot time: N/A
        Next snapshot time: Sep  4 2014 10:38:00
        Next replicable snapshot time: N/A
        Number of snapshots to retain: 30
        Snapshot verification: N/A
        Skip database consistency check: N/A
        Disable appsync: No
        Number of snapshots to retain on replica: 10
        Replicate every: 1
        Latest collection replicated: none
        Replicate from: array7
        Replication in-progress: none
        Created: Feb 19 2014 15:30:50
        Last configuration change: Aug 26 2014 13:28:33
        Last timing configuration change: Mar 5 2014 20:03:20
```

This example creates a new volume collection named *forreplication*. You must also add schedules to the volume collection by using volcoll --addsched.

```
Nimble OS $ volcoll --create forreplication --description "use with
replication" --app_sync none
```

This example modifies the synchronization for a volume collection named *vmwarevollcoll*.

```
Nimble OS $ volcoll --edit vmwarevolcoll --app_sync vmware
--vcenter_hostname 10.12.128.221 --vcenter_username admin
--vcenter_password admin123
```

This example deletes a volume collection named *forreplication*.

```
Nimble OS $ volcoll --delete forreplication
```

This example validates a volume collection named *volcollone*, which is configured with app_sync set to *vss*. You must configure VSS application synchronization before you can validate synchronization using this command.

```
Nimble OS $ volcoll --validate volcollone
```

This example adds a schedule to a volume collection named *forreplication*.

```
Nimble OS $ volcoll --addsched forreplication  --schedule hourly
--repeat 1 --repeat_unit hours --retain 25  --snap_verify no
```

This example edits a schedule for a volume collection named *forreplication*.

```
Nimble OS $ volcoll --editsched forreplication
--schedule hourly --newname everyotherhour
--repeat 2 --repeat_unit hours
```

This example deletes a schedule named *everyotherhour* from a volume collection named *forreplication*.

```
Nimble OS $ volcoll --deletesched forreplication
--schedule everyotherhour
```

This example takes a snapshot collection of a volume collection named *forreplication*.

```
Nimble OS $ volcoll --snap forreplication
--snapcoll_name snap4repl --allow_writes
```

This example promotes the volume collection named *forreplication*.

```
Nimble OS $ volcoll --promote forreplication
```

This example demotes a volume collection named *forreplication*.

```
Nimble OS $ volcoll --demote forreplication
--partner array2
```

This example hands over a volume collection named forreplication to a replication partner named array2 and sets the configuration not to reverse

```
Nimble OS $ volcoll --handover forreplication
--partner array2 --no_reverse
```

This example stops replication for a schedule named *everyotherhour* on a volume collection named *forreplication*.

```
Nimble OS $ volcoll --stop_repl forreplication --schedule everyotherhour
```

This example shows creating a volume collection with app sync type generic:

```
Nimble OS $ volcoll --create gen-sync-test-vc1 --app_sync generic
--agent_hostname 10.xx.xxx.xxx --agent_username Administrator --agent_password
 Nim123#
```

This example shows creating a volume collection from a protection template:

```
Nimble OS $ volcoll --create gen-sync-test-vc2 --prottmpl gen-sync-test-prottmpl
```

This example shows editing volume collection while providing app sync type generic and its arguments:

```
Nimble OS $ volcoll --edit gen-sync-test-vc1 --agent_hostname 10.xx.xxx.xxx
--agent_username admin
      --agent_password Nim123
```

This example shows creating volume collection with an app id of hyperv:

```
Nimble OS $ volcoll --create MyVolColl --app_sync
      vss --app_server 10.xx.xx.xxx --app_id hyperv
```

witness

The witness command enables you to set up and manage the witness host configuration.

Synopsis

witness --help

witness --add --user *username*
--password *password*
--host {*hostname|ipaddr1|ipaddr2|ipaddr..*}
--port *rest_portnum*

witness --info

witness --remove

witness --test

Options

Option	Arguments and Suboptions	Description	Role
--help	N/A	Show usage information for the command.	Administrator
--add	--user *username*	Login username of the witness. This must be a non-root that can login to the witness host (mandatory).	Administrator
	–password *password*	Login password of the new witness.	
	--host {*hostname|ipaddr1,ipaddr2,...*}	Host name or IP addresses of witness (mandatory).	
	--port *rest_portnum*	Port number of witness. If not specified, default port is 5395.	
--info	N/A	List detailed information about the witness configuration.	Administrator
--remove	N/A	Remove an existing witness configuration.	Administrator
--test	N/A	Test the witness configuration between the array and the witness.	Administrator

Regulatory Information

Regulatory Warnings

European Union (EU), Australia, New Zealand

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Israel

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

אזהרה

מוצר זה הוא מוצר Class A.
בסביבה ביתית,מוצר זה עלול לגרום הפרעות בתדר רדיו,ובמקרה זה ,המשתמש עשוי להידרש לנקוט אמצעים מתאימים.

Korea

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

경고:
이 기기는 업무용(A급) 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

Taiwan

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

警告:
警告使用者：這是甲類的資訊產品，在 居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求 採取某 些適當的對策

Thailand

(Compliance Certification Only)

This telecommunication equipment is in compliance with NBTC requirements. (Applies to CS210, CS215, CS235, CS300, CS500, CS700 model arrays, including the ES1-Series expansion shelf.) เครื่องโทรคมนาคม และอุปกรณ์นี้ มีความสอดคล้องตามข้อกำหนดของ กทช.

Saudi Arabia

Nimble Storage product CS200, CS300, CS500, CS700 arrays, and ES1 expansion shelf models comply with the CITC Technical Specification CITC IT001.

USA

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user is required to correct the interference at his own expense.

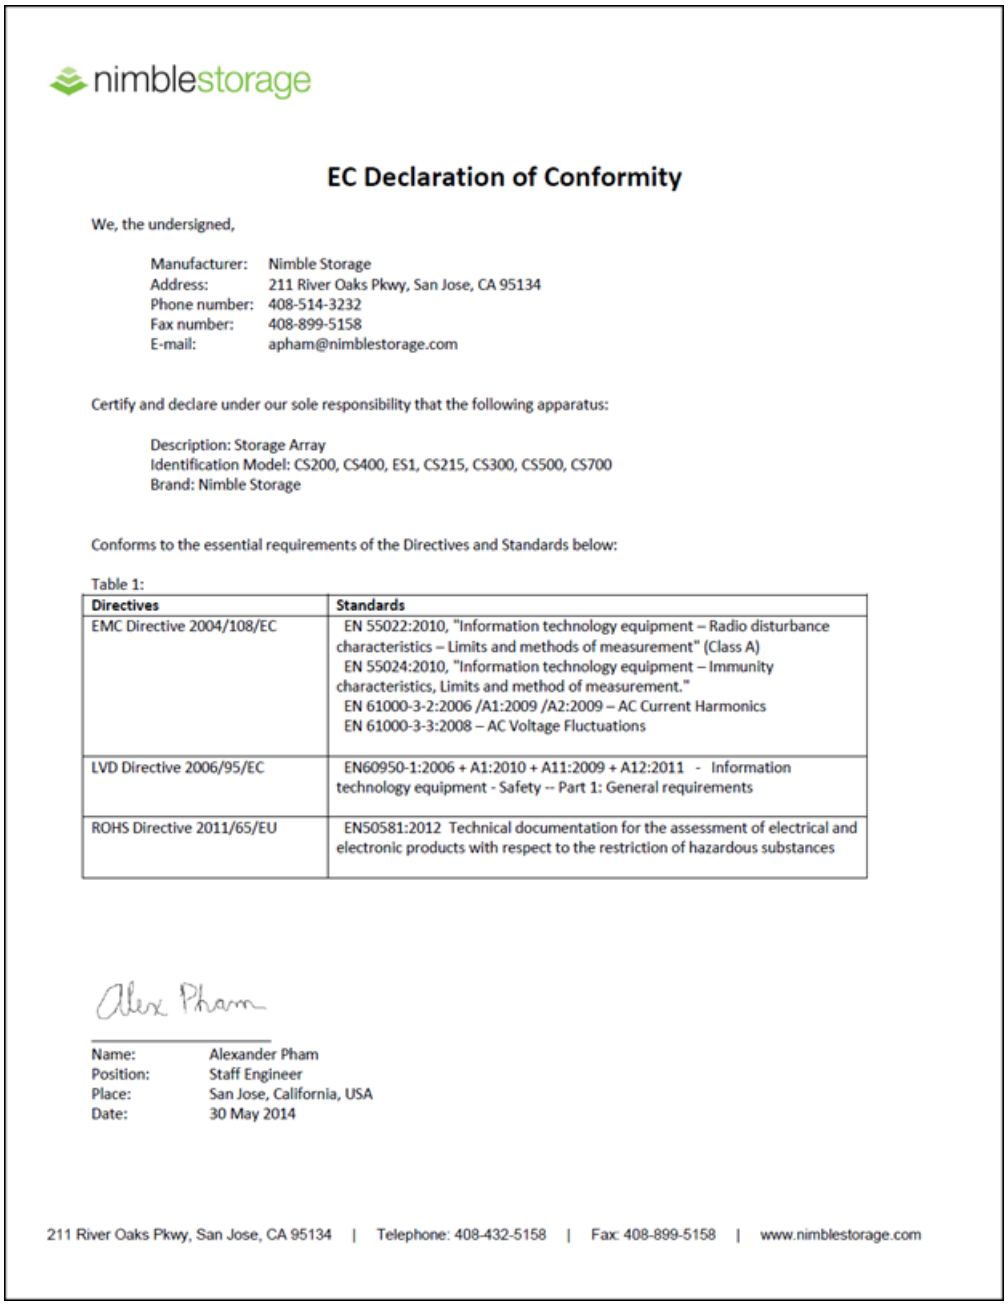

nimblestorage

EC Declaration of Conformity

We, the undersigned,

Manufacturer: Nimble Storage
Address: 211 River Oaks Pkwy, San Jose, CA 95134
Phone number: 408-514-3232
Fax number: 408-899-5158
E-mail: apham@nimblestorage.com

Certify and declare under our sole responsibility that the following apparatus:

Description: Storage Array
Identification Model: CS200, CS400, ES1, CS215, CS300, CS500, CS700
Brand: Nimble Storage

Conforms to the essential requirements of the Directives and Standards below:

Table 1:

Directives	Standards
EMC Directive 2004/108/EC	EN 55022:2010, "Information technology equipment – Radio disturbance characteristics – Limits and methods of measurement" (Class A) EN 55024:2010, "Information technology equipment – Immunity characteristics, Limits and method of measurement." EN 61000-3-2:2006 /A1:2009 /A2:2009 – AC Current Harmonics EN 61000-3-3:2008 – AC Voltage Fluctuations
LVD Directive 2006/95/EC	EN60950-1:2006 + A1:2010 + A11:2009 + A12:2011 - Information technology equipment - Safety -- Part 1: General requirements
ROHS Directive 2011/65/EU	EN50581:2012 Technical documentation for the assessment of electrical and electronic products with respect to the restriction of hazardous substances

Alex Pham

Name: Alexander Pham
Position: Staff Engineer
Place: San Jose, California, USA
Date: 30 May 2014

211 River Oaks Pkwy, San Jose, CA 95134 | Telephone: 408-432-5158 | Fax: 408-899-5158 | www.nimblestorage.com

nimblestorage

Appendix - Declaration of Conformity in Languages of the European Union

- 2004/108/EC (EMC Directive): EN 55022:2010, EN 55024:2010, EN 61000-3-2:2006 /A1:2009 /A2:2009, EN 61000-3-3:2008
- 2006/95/EC (LVD Directive): EN60950-1:2006 + A1:2010 + A11:2009 + A12:2011
- 2011/65/EU (ROHS Directive): EN50581:2012

English	EC Conformity Declaration, in terms of the Guidelines of the European Union: The devices stated below have been developed, constructed and manufactured according to the above mentioned EC directives.
Български (Bulgarian)	ЕО-Декларация за съответствие, по смисъла на Директивите на ЕО: Посочените по-долу продукти са разработени, конструирани и произведени в съответствие с упоменатите по-горе директиви на ЕО.
Čeština (Czech)	ES prohlášení o shodě, ve smyslu směrnic ES: Níže uvedené produkty byly vyvinuty, zkonstruovány a vyrobeny v souladu s výše uvedenými směrnicemi ES.
Dansk (Danish)	EC overensstemmelseerklæring i hensyn til Den Europæiske Unions retningslinjer: Enhederne er angivet nedenfor er udviklet, konstrueret og produceret i overensstemmelse med ovennævnte EU-direktiver.
Nederlands (Dutch)	CE-markering van overeenstemming, conform de EG-richtlijnen: De hieronder vermelde producten zijn ontwikkeld, geconstrueerd en geproduceerd conform de bovengenoemde EG-richtlijnen.
Eesti (Estonian)	EÜ vastavusdeklaratsioon Euroopa Ühenduse suuniste kohaselt. Allnimetatud seade on välja töötatud, konstrueeritud ja toodetud vastavalt ülalnimetatud EÜ direktiividele.
Suomi (Finnish)	Euroopan yhteisön säädöksiin liittyvä EY-vaatimustenmukaisuusvakuutus: alla kuvatut laitteet on kehitetty, kokoonpantu ja valmistettu edellä mainittujen EY-direktiivien mukaisesti.
Français (French)	Déclaration CE de conformité, aux directives de la Communauté européenne: Les appareils présentés ci-dessous ont été développés, construits et conçus conformément aux directives CE susmentionnées.
Deutsch (German)	EG-Konformitätserklärung, im Sinne der EG-Richtlinien: Die unten aufgeführten Produkte wurden entwickelt, konstruiert und gefertigt in Übereinstimmung mit den o.g. EG-Richtlinien.
Ελληνικά (Greek)	Δήλωση συμμόρφωσης ΕΚ, κατά την έννοια των οδηγιών της ΕΚ: Τα προϊόντα που αναφέρονται παρακάτω σχεδιάστηκαν, αναπτύχθηκαν και κατασκευάστηκαν σύμφωνα με τις προαναφερθείσες οδηγίες της ΕΚ.
Magyar (Hungarian)	EK megfelelőségi nyilatkozat az Európai Közösség iránymutatása értelmében: Az alább felsorolt készülékek fejlesztése, tervezése és gyártása a fent említett EK-irányelveknek megfelelően történt.
Italiano (Italian)	Dichiarazione di conformità CE, ai sensi delle direttive CE: I prodotti sotto elencati sono stati sviluppati, costruiti e fabbricati in conformità con le direttive CE sopra indicate.
Latviešu (Latvian)	EK Atbilstības deklarācija, Eiropas kopienas vadlīniju ziņā: Turpmāk minētās ierīces ir izstrādātas, konstruētas un ražotas saskaņā ar iepriekš minētajām EK direktīvām.
Lietuvių (Lithuanian)	EB atitikties deklaracija, atsižvelgiant į Europos bendrijos gaires: Toliau nurodyti prietaisai buvo sukurti, sudedami ir gaminami pagal minetas ES direktyvas.
Norsk (Norwegian)	EC samsvarserklæring, i hensyn til den Europeiske unionens retningslinjer: Enhetene i tabellen nedenfor er utviklet, konstruert og produsert i henhold til de ovenfor nevnte EU-direktiver.
Polski (Poland)	Deklaracja zgodności WE, zgodnie z dyrektywami WE: Określone poniżej urządzenia zostały zaprojektowane, skonstruowane i wyprodukowane zgodnie z powyższymi dyrektywami WE.
Português (Portuguese)	Declaração de Conformidade CE, Conforme as diretivas CE: Os produtos mencionados abaixo foram desenvolvidos, construídos e fabricados em conformidade com as directivas CE acima referidas.
Român (Romanian)	Declarație de conformitate CE, în sensul directivelor CE: Produsele următoare au fost proiectate, construite și fabricate în conformitate cu directivele CE menționate mai sus.
Slovenský (Slovak)	ES Vyhlásenie o zhode, v zmysle smerníc Európskeho spoločenstva: Nižšie uvedené prístroje boli vyvinuté, skonštruované a vyrobené v súlade s vyššie uvedenými smernicami ES.
Slovenščina (Slovenian)	Izjava ES o skladnosti, v smislu direktiv ES: Spodaj navedeni izdelki so bili razviti, zasnovani in izdelani v skladu z zgornjimi direktivami ES.
Svensk (Swedish)	EG-försäkran om överensstämmelse avseende Europeiska gemenskapens riktlinjer: Enheterna som redovisas nedan har utvecklats, konstruerats och tillverkats i enlighet med ovan nämnda EG-direktiv.
Español (Spanish)	Declaración de conformidad de la CE, según las directivas de la CE: Los productos listados abajo fueron desarrollados, diseñados y fabricados conforme a las directivas de la CE arriba indicadas.
Türkçe (Turkish)	AT Uygunluk Beyanı, Avrupa Topluluğu Kuralları açısından: Aşağıda belirtilen cihazlar, yukarıda belirtilen AT yönetmeliklerine göre geliştirilmiş, tasarlanmış ve üretilmiştir.

- Nimble Storage CS200, CS400, ES1, CS215, CS300, CS500, CS700

211 River Oaks Pkwy, San Jose, CA 95134 | Telephone: 408-432-5158 | Fax: 408-899-5158 | www.nimblestorage.com