



**EMC® VPLEX™**  
GeoSynchrony® Release 5.1

**CLI Guide**  
P/N 300-013-927-01

**EMC Corporation**  
*Corporate Headquarters:*  
Hopkinton, MA 01748-9103  
1-508-435-1000  
[www.EMC.com](http://www.EMC.com)

Copyright © May, 2012 EMC Corporation. All rights reserved.

Published May 1, 2012 4:32 pm

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink®.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

<b>Chapter 1</b>	<b>Using the VPLEX CLI</b>	
	Log in to/log out from the CLI.....	10
	CLI context tree .....	11
	Navigate the CLI context tree .....	13
	Using CLI commands.....	17
<b>Chapter 2</b>	<b>VPLEX CLI Command Descriptions</b>	
	advadm dismantle .....	26
	alias .....	27
	array claim .....	29
	array re-discover .....	30
	array used-by.....	32
	authentication directory-service configure .....	34
	authentication directory-service map .....	38
	authentication directory-service show .....	39
	authentication directory-service unconfigure .....	40
	authentication directory-service unmap .....	41
	batch-migrate cancel.....	42
	batch-migrate check-plan .....	43
	batch-migrate clean .....	45
	batch-migrate commit .....	47
	batch-migrate create-plan .....	48
	batch-migrate pause .....	50
	batch-migrate remove .....	51
	batch-migrate resume.....	52
	batch-migrate start.....	53
	batch-migrate summary.....	55
	battery-conditioning disable .....	57
	battery-conditioning enable .....	59
	battery-conditioning manual-cycle cancel-request.....	61
	battery-conditioning manual-cycle request .....	63
	battery-conditioning set-schedule .....	65
	battery-conditioning summary .....	67
	capture begin .....	69
	capture end .....	70
	capture pause .....	71
	capture replay.....	72
	capture resume.....	73
	cd .....	74

chart create .....	75
cluster add .....	77
cluster cacheflush .....	79
cluster configdump .....	81
cluster expel .....	83
cluster forget .....	85
cluster shutdown .....	86
cluster status .....	89
cluster summary .....	93
cluster unexpel .....	95
cluster-witness configure .....	96
cluster-witness disable .....	98
cluster-witness enable .....	101
collect-diagnostics .....	106
configuration complete-system-setup .....	109
configuration configure-auth-service .....	110
configuration connect-local-directors .....	112
configuration connect-remote-directors .....	113
configuration continue-system-setup .....	114
configuration cw-vpn-configure .....	115
configuration cw-vpn-reset .....	118
configuration enable-front-end-ports .....	120
configuration event-notices-reports config .....	121
configuration event-notices-reports reset .....	124
configuration get-product-type .....	126
configuration join-clusters .....	127
configuration metadata-backup .....	128
configuration register-product .....	130
configuration show-meta-volume-candidates .....	132
configuration sync-time .....	133
configuration system-reset .....	134
configuration system-setup .....	135
connect .....	137
connectivity director .....	139
connectivity show .....	141
connectivity validate-be .....	142
connectivity validate-wan-com .....	145
connectivity window set .....	147
connectivity window show .....	149
connectivity window stat .....	150
consistency-group add-virtual-volumes .....	151
consistency-group choose-winner .....	153
consistency-group create .....	156
consistency-group destroy .....	159
consistency-group list-eligible-virtual-volumes .....	160
consistency-group remove-virtual-volumes .....	161
consistency-group resolve-conflicting-detach .....	163
consistency-group resume-after-data-loss-failure .....	166
consistency-group resume-after-rollback .....	169
consistency-group resume-at-loser .....	172
consistency-group set-detach-rule active-cluster-wins .....	174
consistency-group set-detach-rule no-automatic-winner .....	176
consistency-group set-detach-rule winner .....	177
consistency-group summary .....	178
date .....	179

describe.....	180
device attach-mirror .....	181
device collapse.....	183
device detach-mirror .....	184
device resume-link-down .....	186
device resume-link-up .....	188
director appcon .....	189
director appdump.....	190
director appstatus .....	192
director commission .....	193
director decommission.....	194
director fc-port-stats .....	195
director firmware show-banks.....	197
director forget.....	198
director passwd.....	199
director ping .....	200
director shutdown .....	201
director tracepath.....	202
director uptime.....	204
dirs .....	205
disconnect .....	206
dm migration cancel.....	207
dm migration clean.....	208
dm migration commit .....	209
dm migration pause .....	210
dm migration remove .....	211
dm migration resume.....	212
dm migration start.....	213
drill-down .....	217
ds dd create.....	219
ds dd declare-winner .....	222
ds dd destroy .....	223
ds dd remove-all-rules .....	224
ds dd set-log .....	225
ds rule destroy.....	227
ds rule island-containing .....	228
ds rule island-size .....	229
ds rule-set copy .....	230
ds rule-set create .....	231
ds rule-set destroy .....	232
ds rule-set what-if.....	233
ds summary .....	234
event-test .....	239
exec.....	241
exit.....	242
export initiator-port discovery.....	243
export initiator-port register .....	244
export initiator-port register-host.....	246
export initiator-port unregister.....	247
export port summary .....	248
export storage-view addinitiatorport .....	250
export storage-view addport.....	251
export storage-view addvirtualvolume.....	252
export storage-view checkconfig.....	254
export storage-view create .....	255

export storage-view destroy .....	256
export storage-view find .....	257
export storage-view find-unmapped-volumes .....	258
export storage-view map .....	259
export storage-view removeinitiatorport .....	260
export storage-view removeport.....	261
export storage-view removevirtualvolume.....	262
export storage-view show-powerpath-interfaces.....	263
export storage-view summary .....	264
export target-port renamewwns .....	266
extent create .....	267
extent destroy.....	269
extent summary .....	270
getsysinfo.....	272
health-check .....	273
help .....	276
history .....	278
local-device create .....	279
local-device destroy .....	282
local-device summary.....	283
log filter create .....	285
log filter destroy .....	287
log filter list .....	288
log source create .....	289
log source destroy .....	290
log source list .....	291
logging-volume add-mirror.....	292
logging-volume create.....	293
logging-volume destroy .....	295
ls.....	296
management-server set-ip.....	299
manifest upgrade.....	301
manifest version .....	302
meta-volume attach-mirror.....	303
meta-volume backup .....	304
meta-volume create.....	306
meta-volume destroy .....	310
meta-volume detach-mirror .....	311
meta-volume move .....	312
meta-volume verify-on-disk-consistency .....	313
monitor add-console-sink .....	315
monitor add-file-sink.....	317
monitor collect .....	319
monitor create .....	320
monitor destroy .....	323
monitor remove-sink .....	324
monitor stat-list.....	325
ndu pre-config-upgrade.....	327
ndu pre-check .....	330
ndu recover .....	335
ndu rolling-upgrade c41x-to-sles.....	337
ndu rolling-upgrade ssd-fw.....	341
ndu start.....	348
ndu status .....	357
notifications call-home import-event-modifications.....	359

notifications call-home remove-event-modifications .....	361
notifications call-home view-event-modifications.....	363
notifications call-home test.....	365
notifications snmp-trap create .....	366
notifications snmp-trap destroy .....	367
plugin addurl.....	368
plugin listurl .....	369
plugin register .....	370
popd.....	371
ptov describe-be-zoning .....	372
ptov lun-mask-be-storage.....	373
ptov query-be-storage .....	374
ptov suggest-be-zoning-and-masking .....	376
ptov suggest-ports-to-use .....	377
ptov verify-be-storage .....	378
ptov verify-be-zoning.....	379
pushd .....	380
rebuild set-transfer-size .....	381
rebuild show-transfer-size.....	382
rebuild status .....	383
remote-clusters add-addresses .....	384
remote-clusters clear-addresses.....	385
report aggregate-monitors.....	386
report capacity-arrays .....	387
report capacity-clusters.....	390
report capacity-hosts .....	391
report create-monitors.....	392
report poll-monitors .....	395
rp import-certificate.....	396
rp rpa-cluster add .....	397
rp rpa-cluster remove.....	402
rp summary .....	403
rp validate-configuration.....	405
schedule add.....	409
schedule list .....	410
schedule modify .....	411
schedule remove .....	412
scheduleSYR add .....	413
scheduleSYR list.....	414
scheduleSYR remove.....	415
script .....	416
security create-ca-certificate .....	417
security create-certificate-subject.....	421
security create-host-certificate .....	423
security delete-ca-certificate .....	425
security delete-host-certificate .....	426
security export-ca-certificate .....	427
security export-host-certificate .....	428
security import-ca-certificate .....	429
security import-host-certificate.....	431
security ipsec-configure .....	433
security remove-login-banner.....	435
security renew-all-certificates .....	436
security set-login-banner .....	441
security show-cert-subj .....	442

sessions .....	443
set.....	444
set topology .....	447
sms dump .....	448
snmp-agent configure.....	450
snmp-agent start.....	451
snmp-agent status .....	452
snmp-agent stop .....	453
snmp-agent unconfigure .....	454
source .....	455
storage-volume auto-unbanish-interval .....	456
storage-volume claim .....	458
storage-volume claimingwizard .....	461
storage-volume find-array .....	464
storage-volume forget.....	465
storage-volume list-banished .....	466
storage-volume resurrect .....	467
storage-volume summary .....	469
storage-volume unbanish.....	472
storage-volume unclaim.....	474
storage-volume used-by.....	475
subnet clear .....	476
subnet create .....	478
subnet destroy.....	480
subnet modify .....	481
syrcollect.....	482
tree .....	484
unalias .....	485
user add .....	486
user list.....	488
user passwd.....	489
user remove.....	490
user reset.....	491
validate-system-configuration .....	492
vault go .....	493
vault overrideUnvaultQuorum.....	494
vault status .....	496
verify fibre-channel-switches .....	500
version.....	501
virtual-volume create.....	504
virtual-volume destroy.....	507
virtual-volume expand.....	508
virtual-volume summary .....	510
vpn restart .....	512
vpn start.....	513
vpn status .....	514
vpn stop .....	515

## Chapter 3 VPLEX Commands by Topic

### Index

	<b>Title</b>	<b>Page</b>
1	Document Change History .....	3
2	authentication directory-service show field descriptions .....	39
3	batch migration summary field descriptions .....	55
4	battery conditioning field descriptions .....	66
5	battery conditioning summary field descriptions .....	67
6	cluster status field descriptions .....	91
7	cluster summary field descriptions .....	94
8	cluster witness display fields.....	103
9	connectivity window show field descriptions .....	149
10	connectivity window stat field descriptions .....	150
11	Consistency group detach rules .....	174
12	director fc-port-stats field descriptions .....	196
13	director firmware show-banks field descriptions .....	197
14	ds summary field descriptions .....	236
15	extent summary field descriptions .....	271
16	getsysinfo field descriptions .....	272
17	local device summary field descriptions .....	283
18	meta-volume display fields .....	308
19	recoverpoint display fields .....	399
20	rp summary display fields.....	403
21	Create hints files for storage-volume naming.....	462
22	storage-volume summary field descriptions .....	470
23	Vault state field descriptions .....	497
24	Software components .....	503
25	virtual-volume field descriptions .....	505
26	virtual-volume summary field descriptions .....	511



*As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all revisions of the software or hardware currently in use. Your product release notes provide the most up-to-date information on product features.*

*If a product does not function properly or does not function as described in this document, please contact your EMC representative.*

**About this guide**

This document is part of the VPLEX documentation set, and describes the VPLEXcli. Each major section includes introductory information and a general procedure for completing a task.

This guide is part of the VPLEX documentation set, and is intended for use by customers and service providers who configure and manage a storage environment.

**Table 1 Document Change History (1 of 4)**

GeoSynchrony Release	Changes since previous GeoSynchrony release
5.1	<p><b>General</b> - Moved procedural chapters to new <i>EMC VPLEX Administrator's Guide</i>. Chapters moved: Meta-volumes, Distributed Devices, Consistency Groups, and Network Configuration. Moved procedural sections from CLI Overview chapter to <i>EMC VPLEX Administrator's Guide</i>, including add/modify/delete/reset user accounts, and managing the CLI workspace. Removed VPLEX Element Manager API chapter.</p> <p><b>Chapter 1</b> - New contexts in CLI context tree: RecoverPoint (rp) and stand-by-power-supply - new subcontext: conditioning</p> <p><b>Chapter 2:</b>  <u>New commands:</u>                      array used-by, battery-conditioning disable, battery-conditioning enable, battery-conditioning manual-cycle cancel-request, battery-conditioning manual-cycle request, battery-conditioning set-schedule, battery-conditioning summary, configuration event-notices-reports config, configuration event-notices-reports reset, configuration connect-local-directors, director fc-port-stats, notifications call-home import-event-modifications, notifications call-home remove-event-modifications, notifications call-home view-event-modifications, rp import-certificate, rp rpa-cluster add, rp rpa-cluster remove, rp summary, rp validate-configuration, security remove-login-banner, security renew-all-certificates, security set-login-banner</p>

**Table 1 Document Change History (2 of 4)**Modified commands:

batch-migrate start - New arguments: **--force** and **--pause**, increased minimum transfer-size to 40K, changed default transfer-size to 128K

cluster shutdown - New warning messages displayed during execution.

cluster status - New fields to display status of local-com and wan-com

collect-diagnostics - Default behavior changed to collect only the 2 latest core files from directors and to collect diagnostic data for the RecoverPoint splitter. New arguments: **--allcores** to specify collection of all available core files, and **--recoverpoint-only** to collect only RecoverPoint diagnostics.

connectivity validate-be - new output to display storage volumes with too many active paths per director

consistency-group remove-virtual-volumes - changed behavior of **--force** argument

consistency-group create - increased minimum transfer-size to 40K

device attach-mirror - New argument: **--force**.

dm migration start - New arguments: **--force** and **--pause**, Increased minimum transfer-size to 40K, changed default transfer-size to 128K

ds summary - additional output.

export initiator-port register - New type of initiator port: recoverpoint

health-check - New argument: **--full**

management-server set-ip - Modified to support IPv6 addresses for port eth3

ndu rolling-upgrade c41x-to-sles, ndu rolling-upgrade ssd-fw, and ndu start: Removed some skip options

rebuild set-transfer-size - increased minimum limit to 40K

storage-volume claim - modified the meaning of the **--appc** argument

subnet create - increased minimum MTU to 96

Removed commands:

director dns-settings create, director dns-settings destroy<sup>1</sup>

New/modified contexts:

rp - new top-level context for RecoverPoint

stand-by-power-supply - New subcontext: conditioning

5.0.1

**Chapter 1** - Update to illustration of CLI context tree.

**Chapter 2** -

New commands: storage-volume auto-unbanish-interval, storage-volume list-banished, storage-volume unbanish, vault go, vault overrideUnvaultQuorum, vault status

Modified commands:

authentication directory-service configure - new argument: **--custom-attributes**;  
 collect-diagnostics - new argument: **--large-config**; connect - updated description;  
 connectivity validate-wan-com - updated description; consistency-group choose-winner - updated description;  
 consistency-group set-detach-rule active-cluster-wins - updated description;  
 dm migration commands - removed migrations using consistency groups.;  
 local-device create - new arguments : source-leg, **--force**; ndu rolling-upgrade c41x-to-sles - new pre-check skip options;  
 ndu rolling-upgrade ssd-fw - new pre-check skip options; ndu start - new arguments: **--cws-package** and **--force-with-unreachable-cws**, and new pre-check skip options;  
 report create-monitors - new argument: **--force**, and warning if average polling period is too slow;  
 security ipsec-configure - removed two arguments; security create-host-certificate - new argument: **--get-master-ca**;  
 storage-volume resurrect - moved overview text to *EMC VPLEX Product Guide*;

Removed commands:

director coredump add-path, director coredump create, director coredump destroy, director coredump remove-path, static-route create, static-route destroy

**Chapter 4** - Moved overview text to *EMC VPLEX Product Guide*.

**Chapter 6** - Moved overview text to *EMC VPLEX Product Guide*.

**Chapter 9** - Updated list of VPLEX CLI commands supported by VPLEX Element Manager API.

Table 1 Document Change History (3 of 4)

5.0	<p><b>Chapter 1</b> - New context in CLI context tree: consistency groups</p> <p><b>Chapter 2:</b>  <u>New commands:</u>  cluster-witness disable, configuration complete-system-setup, configuration cw-vpn-configure, configuration get-product-type, connectivity validate-wan-com, consistency-group add-virtual-volumes, consistency-group choose-winner, consistency-group create, consistency-group destroy, consistency-group list-eligible-virtual-volumes, consistency-group remove-virtual-volumes, consistency-group resolve-conflicting-detach, consistency-group resume-after-data-loss-failure, consistency-group resume-after-rollback, consistency-group resume-at-loser, consistency-group set-detach-rule active-cluster-wins, consistency-group set-detach-rule no-automatic-winner, consistency-group set-detach-rule winner, consistency-group summary, cluster-witness configure, cluster-witness enable, meta-volume verify-on-disk-consistency, ndu pre-config-upgrade, ndu rolling-upgrade ssd-fw, remote-clusters add-addresses, remote-clusters clear-addresses, security delete-ca-certificate, security delete-host-certificate, subnet clear, subnet create, subnet destroy, subnet modify</p> <p><u>Modified commands:</u>  cluster-witness enable - removed disable option. Replaced with cluster-witness-disable command.  ds dd create - warning when creating a RAID 1 device.  local-device create - --force argument, and warning when creating a RAID 1 device.  storage-volume resurrect: --force option</p> <p><u>Removed commands:</u>  extent claim, extent unclaim, configuration complete-metro-setup (replaced by configuration complete-system-setup)</p> <p><b>Chapter 6</b> - New chapter: Consistency Groups. Overview and procedures for creating and managing synchronous and asynchronous consistency groups.</p> <p><b>Chapter 7</b> - New chapter: Network configuration. Overview and procedure to modify IP configurations.</p> <p><b>Chapter 9</b> - New chapter: VPLEX Element Manager API. Overview and information for creating scripts to configure, monitor and manage VPLEX.</p>
4.2	<p><b>Chapter 1</b> - New section: --verbose and --help arguments.</p> <p><b>Chapter 2</b> - General - reformatted, expanded examples and additional explanation for some commands.</p> <p><u>New commands:</u> array claim, authentication directory-service configure, authentication directory-service map, authentication directory-service show, authentication directory-service unconfigure, authentication directory-service unmap, configuration complete-metro-setup, configuration configure-auth-service, configuration configure-auth-service, configuration join-clusters, configuration metadata-backup, configuration register-product, configuration show-meta-volume-candidates, configuration sync-time, configuration system-reset, configuration sync-time, health-check, local-device summary, ndu rolling-upgrade c41x-to-sles, snmp-agent configure, snmp-agent start, snmp-agent status, snmp-agent stop, snmp-agent unconfigure, user reset, virtual-volume expand</p> <p><u>Removed commands:</u> director bmgr, director firmware activate, director firmware install, director ntp create, director ntp destroy, director qlogic-port-errors, director reboot, director run, director runurl, director stop, director syslog create, director syslog destroy, local-device expand, meta-volume activate, monitor add-snmp-sink, ndu cleanup.</p> <p><u>Commands obsoleted, but not yet removed from the CLI:</u> extent claim, extent unclaim</p> <p><b>Chapter 6</b> - New sections: Pre-configured monitors, Monitoring with SNMP, Performance statistics retrievable by SNMP.</p>

Table 1 Document Change History (4 of 4)

4.1	<p><b>Chapter 1</b> - Reformatted. New figure: VPLEX CLI context tree</p> <p><b>Chapter 2</b> - Reformatted command descriptions to interim format.</p> <p><u>New/changed commands</u>: ndu cleanup, ndu pre-check, ndu recover, ndu start, ndu status, ptov describe-be-zoning, ptov lun-mask-be-storage, ptov query-be-storage, ptov suggest-be-zoning-and-masking, ptov suggest-ports-to-use, ptov verify-be-storage, ptov verify-be-zoning, verify fibre-channel-switches.</p> <p><u>Corrected commands</u>: array re-discover, cd, ds dd create, describe, meta-volume create, user add, user passwd.</p> <p><b>Chapter 3</b> - New chapter: Meta-volumes. Procedures to create, backup, activate, rename, move, delete, and display meta-volumes.</p> <p><b>Chapter 4</b> - New chapter: Data mobility. Procedures to start, monitor, pause, resume, cancel, commit, clean, and remove one-time and batch data migrations.</p> <p><b>Chapter 5</b> - Reformatted. New overview material for logging volumes, rule-sets, and distributed devices.</p> <p><b>Chapter 6</b> - New chapter: Performance monitoring. Procedures to create, modify, and delete performance monitors, and the statistics that can be monitored.</p> <p><b>Chapter 7</b> - Reformatted. New tables.</p>
-----	---

### Related documentation

Related documents (available on EMC Powerlink) include:

- ◆ *EMC VPLEX Release Notes for GeoSynchrony Releases 5.0 and 5.1*
- ◆ *EMC VPLEX Product Guide*
- ◆ *EMC VPLEX Administrator's Guide*
- ◆ *EMC VPLEX Site Preparation Guide*
- ◆ *EMC VPLEX Configuration Worksheet*
- ◆ *EMC VPLEX Configuration Guide*
- ◆ *EMC VPLEX Hardware Installation Guide*
- ◆ *Unisphere for VPLEX Help*
- ◆ *EMC VPLEX Security Configuration Guide*
- ◆ *EMC VPLEX Open-Source Licenses*
- ◆ *EMC Regulatory Statement for EMC VPLEX*
- ◆ EMC VPLEX Procedure Generator
- ◆ EMC Host Connectivity Guides

### Conventions used in this document

EMC uses the following conventions for special notices.

Note: A note presents information that is important, but not hazard-related.



**A caution contains information essential to avoid data loss or damage to the system or equipment.**

#### IMPORTANT

**An important notice contains information essential to operation of the software.**

## Typographical conventions

EMC uses the following type style conventions in this document:

<b>Normal</b>	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)</li> <li>Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, utilities</li> <li>URLs, pathnames, filenames, directory names, computer names, filenames, links, groups, service keys, file systems, notifications</li> </ul>
<b>Bold</b>	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>Names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system call, man pages</li> </ul> Used in procedures for: <ul style="list-style-type: none"> <li>Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)</li> <li>What user specifically selects, clicks, presses, or types</li> </ul>
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> <li>Full titles of publications referenced in text</li> <li>Emphasis (for example a new term)</li> <li>Variables</li> </ul>
<i>Courier</i>	Used for: <ul style="list-style-type: none"> <li>System output, such as an error message or script</li> <li>URLs, complete paths, filenames, prompts, and syntax when shown outside of running text</li> </ul>
<b>Courier bold</b>	Used for: <ul style="list-style-type: none"> <li>Specific user input (such as commands)</li> </ul>
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> <li>Variables on command line</li> <li>User input variables</li> </ul>
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

## Where to get help

EMC support, product, and licensing information can be obtained as follows.

**Product information** — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Powerlink website (registration required) at:

<http://Powerlink.EMC.com>

**Technical support** — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

## Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this document to:

[techpubcomments@EMC.com](mailto:techpubcomments@EMC.com)



This chapter describes how to use the VPLEX™ command line interface (CLI).

- ◆ Log in to/log out from the CLI ..... 14
- ◆ CLI context tree..... 15
- ◆ Navigate the CLI context tree..... 17
  
- ◆ Using CLI commands ..... 21
- ◆ Tab completion ..... 22
- ◆ Wildcards ..... 23
- ◆ Names ..... 24
- ◆ Command globbing ..... 24
- ◆ --verbose and --help arguments ..... 25
- ◆ Get help..... 26

## Log in to/log out from the CLI

### Log in

1. Log in to the service account at the cluster's management server's public IP address using an SSH client (PuTTY or OpenSSH). Configure the SSH client as follows:

- Port 22
- SSH protocol version is set to 2
- Scrollback lines set to 20000

The log in prompt appears:

```
login as:
```

2. Type **service** and press **ENTER**. A password prompt appears:

```
Using keyboard-interactive authentication.
```

```
Password:
```

3. Type the service password and press **ENTER**. The default password is **Mi@Dim7T**. A server prompt appears:

```
service@ManagementServer:~>
```

4. Type the **vplexcli** command to connect to the VPLEXcli:

```
service@ManagementServer:~> vplexcli
```

Several messages are displayed, and a username prompt appears:

```
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^['.
```

```
Enter User Name:
```

5. Type **service** and press **ENTER**.

```
Enter User Name: service
```

A password prompt appears:

```
Password:
```

6. Type the service password and press **ENTER**. The default password is **Mi@Dim7T**.

The VPLEXcli prompt appears:

```
creating logfile:/var/log/Vplex/cli/session.
log_service_localhost_T28921_20101020175912
```

```
Vplexcli: />
```

### Log out

Use the **exit** command to exit the VPLEXcli from any context.

For example:

```
Vplexcli:/clusters> exit
Connection closed by foreign host.
```

## CLI context tree

The CLI is divided into command contexts. Some commands are accessible from all contexts, and are referred to as 'global commands'.

The remaining commands are arranged in a hierarchical context tree. These commands can only be executed from the appropriate location in the context tree.

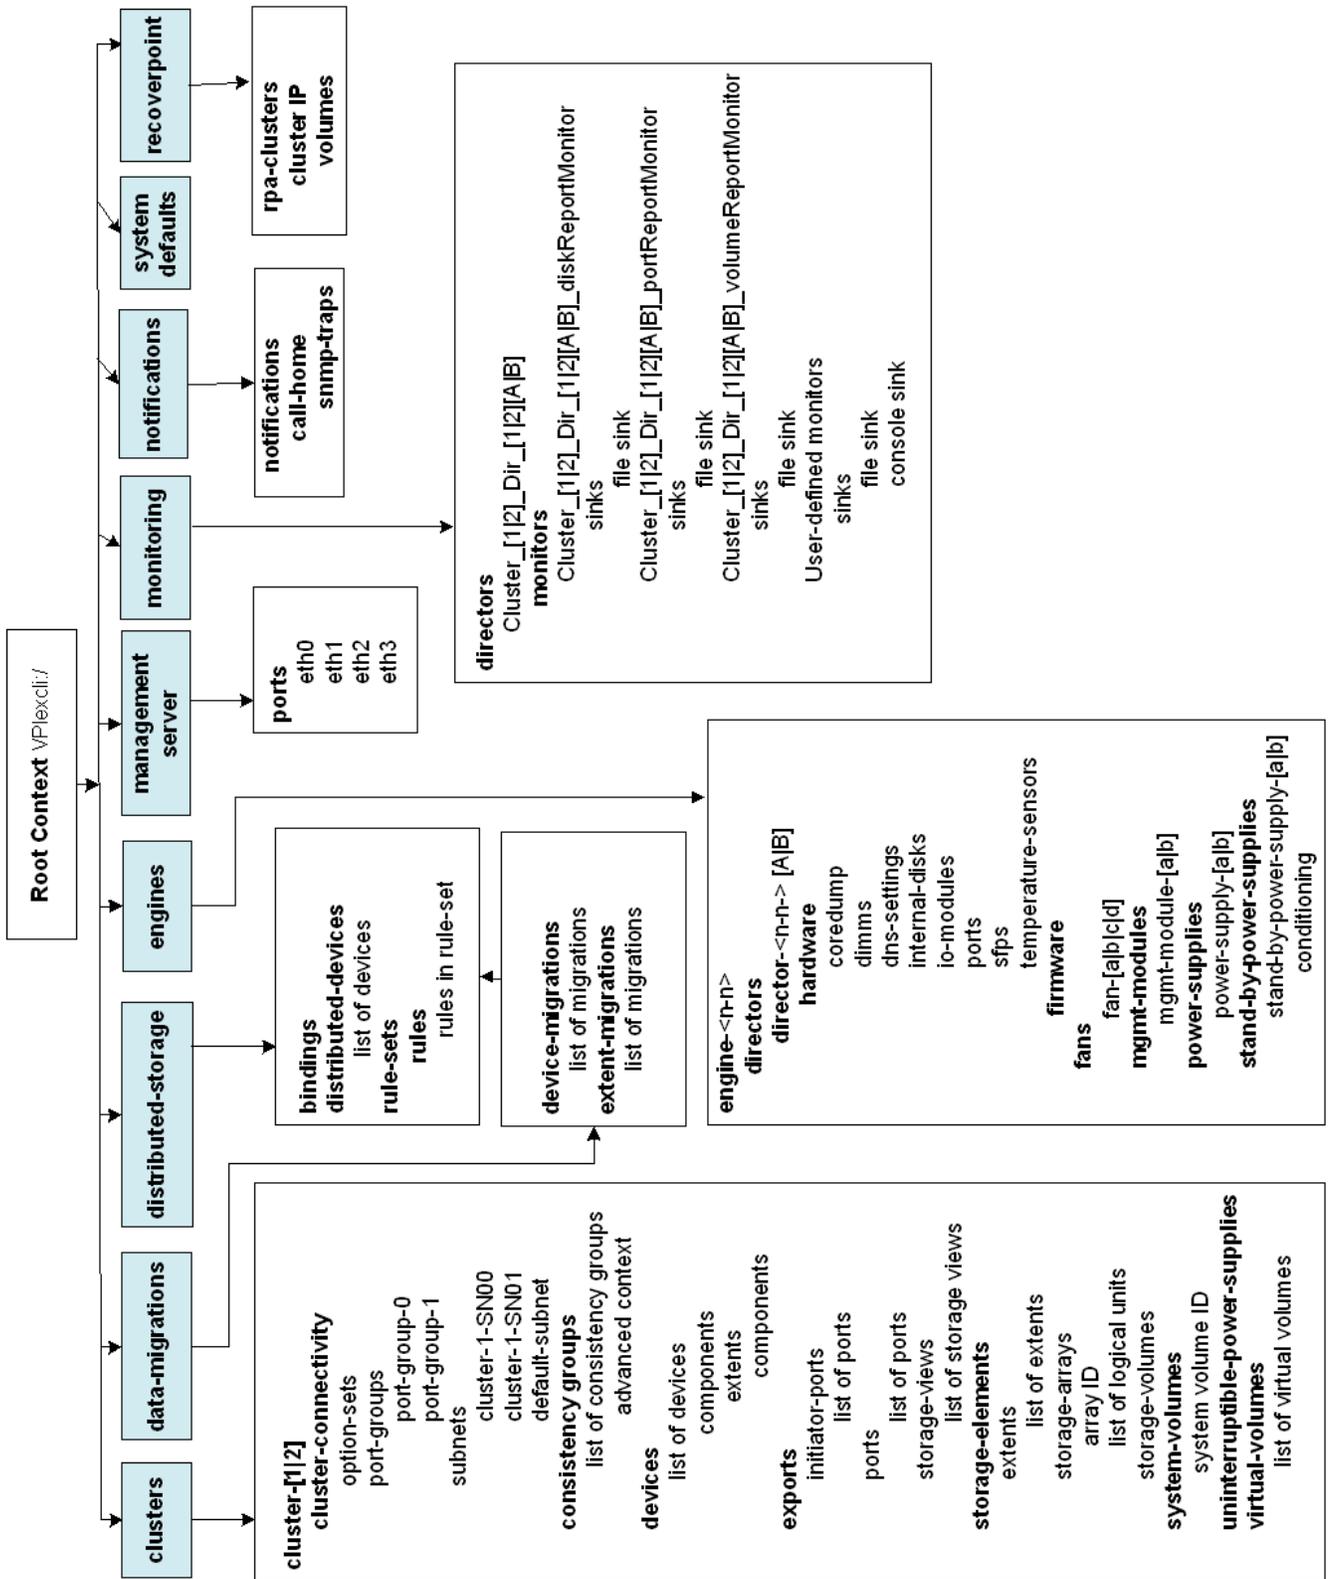
Understanding the command context tree is critical to using the VPLEXcli effectively.

The root context contains eight sub-contexts:

- ◆ **clusters/** - Create and manage links between clusters, devices, extents, system volumes and virtual volumes. Register initiator ports, export target ports, and storage views.
- ◆ **data-migrations/** - Create, verify, start, pause, cancel, and resume data migrations of extents or devices.
- ◆ **distributed-storage/** - Create and manage distributed devices and rule sets.
- ◆ **engines/** - Configure and manage directors, fans, management modules, and power.
- ◆ **management-server/** - Manage the Ethernet ports.
- ◆ **monitoring/** - Create and manage performance monitors.
- ◆ **notifications/** - Create and manage call-home events.
- ◆ **system-defaults/** - Display systems default settings.

Except for **system-defaults/**, each of the eight sub-contexts contains one or more sub-contexts to configure, manage, and display sub-components.

Command contexts have commands that can be executed only from that context. The command contexts are arranged in a hierarchical context tree. The topmost context is the root context, or `"/"`.



## Navigate the CLI context tree

Use the **cd** command to navigate between command contexts.

The current context is always displayed at the VPLEXcli prompt:

```
VPLEXcli:/> cd /clusters/cluster-1/devices/
```

```
VPLEXcli:/clusters/cluster-1/devices>
```

For example, to navigate from the root (/) context to the monitoring context for a specified director:

```
VPLEXcli:/> cd /monitoring
VPLEXcli:/monitoring> cd /directors/
VPLEXcli:/monitoring/directors> cd /director-1-1-B
VPLEXcli:/monitoring/directors/director-1-1-B> cd /monitors
VPLEXcli:/monitoring/directors/director-1-1-B/monitors> cd
/director-1-1-B
VPLEXcli:/monitoring/directors/director-1-1-B/monitors/director-1-1-B
> ls
Attributes:
Name          Value
-----
average-period 1s
bucket-count  64
.
.
.
```

Alternatively, type all the context identifiers in a single command. For example, the above navigation can be typed as:

```
VPLEXcli:/> cd
/monitoring/directors/director-1-1-B/monitors/director-1-1-B_director
VPLEXcli:/monitoring/directors/director-1-1-B/monitors/director-1-1-B
>
```

Use the **cd** command with no arguments or followed by three periods (**cd...**) to return to the root context:

```
VPLEXcli:/engines/engine-1-1/fans> cd
```

```
VPLEXcli:/>
```

Use the **cd** command followed by two periods (**cd..**) to return to the context immediately above the current context:

```
VPLEXcli:/monitoring/directors/director-1-1-B> cd ..
```

```
VPLEXcli:/monitoring/directors>
```

To navigate directly to a context from any other context use the **cd** command and specify the absolute context path. In the following example, the **cd** command changes the context from the data migrations/ extent-migrations context to the engines/engine-1/fans context:

```
VPLEXcli:/data-migrations/extent-migrations> cd
/engines/engine-1-1/fans/
```

```
VPLEXcli:/engines/engine-1-1/fans>
```

### pushd and popd commands

- ◆ Use the **pushd *directory*** command to save the current directory, and jump to the specified directory.

Once a directory is added to the pushd stack, use the **pushd** command with no argument to switch back to the previous directory.

In the following example, **pushd** toggles between the engines and monitoring parent contexts:

```
Vplexcli:/engines/engine-1-1/directors/director-1-1-A> pushd
/monitoring/directors/director-1-1-A
[/monitoring/directors/director-1-1-A,
/engines/engine-1-1/directors/director-1-1-A,
/monitoring/directors/director-1-1-A]
```

```
Vplexcli:/monitoring/directors/director-1-1-A> pushd
[/engines/engine-1-1/directors/director-1-1-A,
/monitoring/directors/director-1-1-A,
/monitoring/directors/director-1-1-A]
```

```
Vplexcli:/engines/engine-1-1/directors/director-1-1-A> pushd
[/monitoring/directors/director-1-1-A,
/engines/engine-1-1/directors/director-1-1-A,
/monitoring/directors/director-1-1-A]
```

```
Vplexcli:/monitoring/directors/director-1-1-A>
```

- ◆ Use the **dirs** command to display to the current context stack:

```
Vplexcli:/clusters/cluster-1> dirs
[/clusters/cluster-1, /, /,
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A5-GE0
1, /]
```

- ◆ Use the **popd** command to remove the last directory saved by the **pushd** command and jump to the new top directory.

In the following example, the **dirs** command displays the context stack saved by the **pushd** command, and the **popd** command removes the top directory, and jumps to the new top directory:

```
Vplexcli:/engines/engine-1-1/directors/director-1-1-A> dirs
[/engines/engine-1-1/directors/director-1-1-A,
/monitoring/directors/director-1-1-A]
```

```
Vplexcli:/engines/engine-1-1/directors/director-1-1-A> popd
[/engines/engine-1-1/directors/director-1-1-A]
```

```
Vplexcli:/monitoring/directors/director-1-1-A>
```

---

**Note:** The context tree displays only those objects associated with directors to which the management system is connected.

---

- ◆

```
Vplexcli: /> cd /monitoring/directors/director-1-1-B/monitors/
```

```
Vplexcli:/monitoring/directors/director-1-1-B/monitors>
```

- ◆ The **ls** command displays the sub-contexts immediately accessible from the current context:

```
Vplexcli: /> ls
clusters          data-migrations   distributed-storage
engines           management-server  monitoring
```

```
notifications system-defaults
```

- ◆ The **ls -l** command displays more information about the current sub-contexts:

```
Vplexcli:/data-migrations> ls -l
Name                Description
-----
device-migrations  Contains all the device migrations in the system.
extent-migrations  Contains all the extent migrations in the system.
```

- ◆ For contexts where the next lowest level is a list of individual objects, the **ls** command displays a list of the objects:

```
Vplexcli:/clusters/cluster-1/exports/ports> ls
P000000003B2017DF-A0-FC00  P000000003B2017DF-A0-FC01
P000000003B2017DF-A0-FC02  P000000003B2017DF-A0-FC03
P000000003B3017DF-B0-FC00  P000000003B3017DF-B0-FC01
P000000003B3017DF-B0-FC02  P000000003B3017DF-B0-FC03
```

- ◆ <Tab>

For example, type **cd** and press <Tab> in the data-migrations context to display available options:

```
Vplexcli:/data-migrations> cd <Tab>
device-migrations/  extent-migrations/
```

- ◆ The **tree** command displays the immediate sub-contexts in the tree using the current context as the root:

```
Vplexcli:/ cd /clusters/cluster-1/devices/Symm_rC_3

Vplexcli:/clusters/cluster-1/devices/Symm_rC_3> tree
/clusters/cluster-1/devices/Symm_rC_3:
  components
    Symm_rC_3_extent_0
    Symm_rC_3_extent_1
```

- ◆ The **tree -e** command displays immediate sub-contexts in the tree and any sub-contexts under them:

```
Vplexcli:/clusters/cluster-1/devices/Symm_rC_3> tree -e
/clusters/cluster-1/devices/Symm_rC_3:
  components
    Symm_rC_3_extent_0
      components
        Symm0487_44C
          components
            Symm_rC_3_extent_1
              components
                Symm0487_44B
                  components
```

---

**Note:**

---

```
Vplexcli:/clusters/cluster-1> tree
/clusters/cluster-1:
  cluster-connectivity
  cluster-links
  to-cluster-2
  proxy-servers
  static-routes
  devices
  base0
  components
  extent_CX4_lun0_1
```

```
        components
        CX4_lun0
        components
    .
    .
    .
exports
    initiator-ports
        LicoJ006_hba0
        LicoJ006_hba1
    .
    .
    .
ports
        P000000003CA00147-A0-FC00
        P000000003CA00147-A0-FC01

storage-views
    LicoJ009
    LicoJ013
storage-elements
    extents
        extent_CX4_Logging_1
```

## Using CLI commands

This section describes the following topics:

◆ Display currently available commands .....	21
◆ Page output .....	22
◆ Tab completion .....	22
◆ Wildcards .....	23
◆ Command globbing .....	24
◆ Positional command arguments .....	25
◆ --verbose and --help arguments .....	25
◆ Search command history .....	26
◆ View command history .....	26
◆ Get help.....	26

### Display currently available commands

The commands that make up the CLI fall into two groups:

- ◆ Global commands that can be used in any context. For example: **cd**, **date**, **ls**, **exit**, **user**, and **security**.
- ◆ Context-specific commands that can be used only in specific contexts. For example, to use the **copy** command, the context must be `/distributed-storage/rule-sets`.

Use the **help** command to display a list of all commands (including the global commands) available from the current context.

Use the **help -G** command to display a list of available commands in the current context *excluding* the global commands:

```
Vplexcli:/notifications> help -G
Commands specific to this context and below:
call-home snmp-trap
```

Some contexts “inherit” commands from their parent context. These commands can be used in both the current context and the context immediately above in the tree:

```
Vplexcli:/distributed-storage/bindings> help -G
Commands inherited from parent contexts:
dd rule rule-set summary
```

Some commands are loosely grouped by function. For example, the commands to create and manage performance monitors start with the word “monitor”.

Use the <Tab> key display the commands within a command group. For example, to display the commands that start with the word “monitor”, type “monitor” followed by the <Tab> key:

```
Vplexcli:/> monitor <Tab>

add-console-sink  add-file-sink      collect           create
destroy           remove-sink
stat-list
```

## Page output

For large configurations, output from some commands can reach hundreds of lines.

Paging displays long output generated by the `ll` and `ls` commands one “page” at a time:

To enable paging, add `-p` at the end of any command:

```
Vplexcli:/clusters/cluster-1/storage-elements> ls storage-volumes -p
```

One page of output is displayed. The following message is at the bottom of the first page:

```
-- more --(TOP )- [h]elp
```

Press the spacebar to display the next page.

The message now indicates what percentage of the output has been displayed:

```
-- more --( 24%)- [h]elp
```

**h** - Displays instructions on how to move and search the output.

**q** - Exits paging mode.

## Tab completion

Use the <Tab> key to:

- ◆ Complete a command
- ◆ Display valid contexts and commands
- ◆ Display command arguments

### Complete a command

Use the <Tab> key to automatically complete a path or command until the path/command is no longer unique.

For example, to navigate to the UPS context on a single cluster (named cluster-1), type:

```
cd /clusters/cluster-1/uninterruptible-power-supplies/
```

To type the same command using tab completion:

1. Type `cd c <Tab>`

Since 'clusters' is the only context starting with 'c' at the root level, the CLI auto-completes the selection:

```
cd /clusters/
```

2. There is only one cluster (it is unique). Press <Tab> to automatically specify the cluster:

```
cd /clusters/cluster-1/
```

3. Type a `u` to select the uninterruptible-power-supplies context and press <Tab>.

The `u` is unique at the current context, and the CLI auto-completes the selection:

```
cd /clusters/cluster-1/uninterruptible-power-supplies/
```

### Display valid contexts and commands

Press <Tab> after typing a partial context path to display a list of valid commands and/or contexts for the current context:

```
Vplexcli:/> cd /clusters/cluster-1/ <Tab>
```

```

cluster-connectivity/
exports/
system-volumes/
virtual-volumes/

devices/
storage-elements/
uninterruptible-power-supplies/

```

```
Vplexcli: /> cd /clusters/cluster-1/
```

### Display command arguments

Press <Tab> after typing a command name to display the command's arguments. For example:

```
Vplexcli: /> monitor <Tab>

add-console-sink    add-file-sink    collect
create              destroy          remove-sink stat-list
```

## Wildcards

The Vplexcli includes 3 wildcards:

- \* - matches any number of characters.
- ? - matches any single character.
- [a|b|c] - matches any of the single characters a or b or c.

### \* wildcard

Use the \* wildcard to apply a single command to multiple objects of the same type (directors or ports). For example, to display the status of ports on each director in a cluster, without using wildcards:

```

ll engines/engine-1-1/directors/director-1-1-A/hardware/ports
ll engines/engine-1-1/directors/director-1-1-B/hardware/ports
ll engines/engine-1-2/directors/director-1-2-A/hardware/ports
ll engines/engine-1-2/directors/director-1-2-B/hardware/ports
.
.
.

```

Alternatively:

- ◆ Use one \* wildcard to specify all engines, and
  - ◆ Use a second \* wildcard specify all directors:
- ```
ll engines/engine-1-*/directors/*/hardware/ports
```

### \*\* wildcard

Use the \*\* wildcard to match all contexts and entities between two specified objects. For example, to display all director ports associated with all engines without using wildcards:

```

ll /engines/engine-1-1/directors/director-1-1-A/hardware/ports
ll /engines/engine-1-1/directors/director-1-1-B/hardware/ports
.
.
.

```

Alternatively, use a \*\* wildcard to specify all contexts and entities between /engines and ports:

```
ll /engines/**/ports
```

### ? wildcard

Use the ? wildcard to match a single character (number or letter).

```
ls /storage-elements/extents/0x1?[8|9]
```

Returns information on multiple extents.

### [a|b|c] wildcard

Use the [a|b|c] wildcard to match one or more characters in the brackets.

```
ll engines/engine-1-1/directors/director-1-1-A/hardware/ports/A[0-1]
```

displays only ports with names starting with an A, and a second character of 0 or 1.

## Names

Major components of the VPLEX are named as follows:

- ◆ Clusters - VPLEX Local™ configurations have a single cluster, with a cluster ID of 1. VPLEX Metro™ and VPLEX Geo™ configurations have two clusters with cluster IDs of 1 and 2.

```
Vplexcli:/clusters/cluster-1/
```

- ◆ Engines are named *<engine-n-n>* where the first value is the cluster ID (1 or 2) and the second value is the engine ID (1-4).

```
Vplexcli:/engines/engine-1-2/
```

- ◆ Directors are named *<director-n-n-n>* where the first value is the cluster ID (1 or 2), the second value is the engine ID (1-4), and the third is A or B.

```
Vplexcli:/engines/engine-1-1/directors/director-1-1-A
```

For objects that can have user-defined names, names must comply with the following rules:

- ◆ Can contain uppercase and lowercase letters, numbers, and underscores
- ◆ No spaces
- ◆ Cannot start with a number
- ◆ No more than 63 characters

## Command globbing

Command “globbing” combines wildcards and context identifies in a single command. Globbing can address multiple entities using a single command.

### Example 1

In the following example, a single command enables ports in all engines and all directors (A and B) whose name include 0-FC and 1-FC:

```
set /engines/*/directors/*/hardware/ports/*[0-1]-FC*:: enabled true
```

- ◆ First \* wildcard — All engines in the cluster.
- ◆ Second \* wildcard — All directors in the cluster.
- ◆ Third \* wildcard — All A-side ports and all B-side ports.
- ◆ The [0-1] limits the selections to all port numbers that start with A0, A1, B0, or B1.
- ◆ Fourth \* wildcard — All ports whose numbers start with A0-FC, A1-FC, B0-FC, or B1-FC.

### Example 2

To display the status of all the director ports on a large configuration using no wildcards, type:

```
ll /engines/engine-1-<Enclosure_ID>/directors/<director_
name>/hardware/ports
```

for each engine and director.

Using the \* wildcard reduces this task to a single command:

```
ll /engines/engine-1-*/directors/*/hardware/ports
```

Using the **\*\*** wildcard simplifies the command even more:

```
ll /**/ports
```

## Positional command arguments

Most commands require arguments.

Some command arguments are *positional*. That is, the argument can be typed without an identifier IF it is entered in the position specified by the command syntax.

For example, the **alias** command has two arguments in the following order (syntax):

```
alias
[-n|--name] <alias name>
[-t|to] <"string of commands in quotes">
```

Type the command with the arguments *with identifiers* in any order (not as specified by the syntax):

```
Vplexcli:/> alias --to "cd clusters" --name cdc
or,
```

Type the command with the arguments *without identifiers* in the order specified by the command syntax:

```
Vplexcli:/> alias cdc "cd clusters"
```

## --verbose and --help arguments

All commands have two arguments:

- ◆ **--help** displays the online help for the specified command. For example:

```
Vplexcli:/> exec --help
synopsis: exec [<options>] <command name> <word>*
```

Executes an external program.

options (\* = required):

```
.
.
.
```

See "Get help" on page 26.

- ◆ **--verbose** displays additional information for some commands. For example, without **--verbose** argument:

```
Vplexcli:/> connectivity validate-be
```

Summary

Cluster cluster-1

This cluster has 0 storage-volumes which do not have dual paths

This cluster has 0 storage-volumes which are not visible from all directors

With **--verbose** argument:

```
Vplexcli:/> connectivity validate-be --verbose
```

Storage volumes that are dead or unreachable:

```
Cluster    Dead or Unreachable Storage Volumes
-----
```

```
cluster-2 VPD83T3:60004530000000080007f16e9512a2b1
cluster-1 VPD83T3:60004530000000010007f16e9512a2a5
          VPD83T3:60004530000000010007f16e9512a2a7
          VPD83T3:60004530000000010007f16e9512a2a9
```

#### Summary

##### Cluster cluster-2

```
This cluster has 1 storage-volumes which are dead or unreachable
This cluster has 0 storage-volumes which do not have dual paths
This cluster has 0 storage-volumes which are not visible from all directors
```

##### Cluster cluster-1

```
This cluster has 3 storage-volumes which are dead or unreachable
This cluster has 0 storage-volumes which do not have dual paths
This cluster has 0 storage-volumes which are not visible from all directors
```

---

## Search command history

- ◆ To display the last commands typed, press the up arrow key.
- ◆ To search for a command typed in the current CLI session, press **Ctrl-r**.

The reverse search prompt is displayed:

```
(reverse-i-search) '':
```

Type the first letter of the command to search for. After the first letter is typed, the search tool displays a list of possible matches.

---

## View command history

Use the “up arrow” key to display the last command typed.

Use the “up arrow” key, multiple times to display recent command history.

Use the **history** command to display a complete list of commands executed in the current session:

```
Vplexcli:/engines/engine-0-0/directors> history
Vplexcli:/> history
0 cd engines/engine-0-0/directors
1 extent unclaim *
2 ls
3 ls -l
4 extent claim *
5 ls
6 ls -l
7 ls -la
```

Use the **history *nn*** command to display the last *nn* entries in the list:

```
Vplexcli:/clusters/cluster-1> history 22
478 ls storage-volumes -p
479 cd clusters/cluster-1/
480 ls storage-volumes
481 cd storage-elements/
482 ls storage-volumes -p
```

---

## Get help

- ◆ Use the **help** or **?** command with no arguments to display all the commands available, the current context, including global commands.
- ◆ Use the **help** or **?** command with **-G** argument to display all the commands available, the current context, excluding global commands:

```
Vplexcli:/clusters> help -G
```

```
Commands specific to this context and below:
```

```
add cacheflush configdump expel forget shutdown summary unexpel
```

- ◆ Use the **help command** or **command --help** command to display help for the specified command.



This chapter describes each of the VPLEX CLI commands in alphabetical order.

## advadm dismantle

Dismantles VPLEX storage objects down to the storage-volume level, and optionally unclaims the storage volumes.

**Contexts** All contexts.

**Syntax**

```
advadm dismantle
  [-r|--devices] context path,context path
  [-v|--virtual-volumes] context path,context path
  [--unclaim-storage-volumes]
  [-f|--force]
```

**Arguments** **Required arguments**

**[-r|--devices] context path,context path...** - One or more device(s) to dismantle. Entries must be separated by commas. Glob patterns may be used.

**[-v|--virtual-volumes] context path,context path...** - One or more virtual volume(s) to dismantle. Entries must be separated by commas. Glob patterns may be used.

**Optional arguments**

**--unclaim-storage-volumes** - Unclaim the storage volumes after the dismantle is completed.

**[-f|--force]** - Force the dismantle without asking for confirmation.

**Description** If the **--force** argument is used, no confirmation is displayed before the dismantle is performed.

**Example**

```
Vplexcli:/clusters/cluster-1> advadm dismantle --verbose
--virtual-volumes virtual-volumes/test_r1_vol --force
```

```
destroyed virtual volume
  /clusters/cluster-1/virtual-volumes/test_r1_vol
destroyed /clusters/cluster-2/devices/test_r1
Destroyed 1 out of 1 targetted extents.

destroyed
  /clusters/cluster-1/storage-elements/extents/extent_CLAR0014_LUN14_
  1
```

**See also**

- ◆ [ds dd create on page 219](#)
- ◆ [local-device create on page 279](#)
- ◆ [virtual-volume create on page 504](#)

## alias

Creates a command alias.

**Contexts** All contexts.

**Syntax**

```
alias
  [-n|--name] name
  [-t|--to] "commands and arguments to include
            in the new alias"
```

**Arguments** **Required arguments**

**[-n | --name] name** - \* The name of the new alias.

- Up to 63 characters.
- May contain letters, numbers, and underscores '\_'.
- Cannot start with a number.

**[-t | --to] "commands and arguments"** - \* A string of commands and arguments enclosed in quote marks. This string is invoked when the aliased command is used.

\* - argument is positional.

**Description** Aliases are shortcuts for frequently used commands or commands that require long strings of context identifiers.

Use the **alias** command with no arguments to display a list of all aliases configured on the VPLEX.

Use the **alias name** command to display the underlying string of the specified alias.

Use the **alias name "string of CLI commands"** command to create an alias with the specified name that invokes the specified string of commands.

Use the **"unalias"** command to delete an alias.

The following aliases are pre-configured:

- **?** Substitutes for the **help** command.
- **ll** Substitutes for the **ls -a** command.
- **quit** Substitutes for the **exit** command.

An alias that executes correctly in one context may conflict with an existing command when executed from another context (pre-existing commands are executed before aliases if the syntax is identical).

Make sure that the alias **name** is unique, that is, not identical to an existing command or alias.

If an alias does not behave as expected, delete or reassign the alias.

The CLI executes the first matching command, including aliases in the following order:

1. Local command in the current context.
2. Global command in the current context.
3. Root context is searched for a match.

**Example** Create an alias:

```
VPlxcli:/> alias mon-Dir-1-1-B "cd /monitoring/directors/director-1-1-B"
```

Display a list of aliases:

```
VPLEXcli:/> alias
```

| Name          | Description                                               |
|---------------|-----------------------------------------------------------|
| ?             | Substitutes the 'help' command.                           |
| mon-Dir-1-1-B | Substitutes the 'cd /monitoring/directors/director-1-1-B' |
| ll            | Substitutes the 'ls -al' command.                         |
| quit          | Substitutes the 'exit' command.                           |

Display a specified alias:

```
VPLEXcli:/> alias mon-Dir-1-1-B
```

| Name          | Description                                                        |
|---------------|--------------------------------------------------------------------|
| mon-Dir-1-1-B | Substitutes the 'cd /monitoring/directors/director-1-1-B' command. |

Use an alias:

```
VPLEXcli:/> mon-Dir-1-1-B
```

```
VPLEXcli:/monitoring/directors/director-1-1-B>
```

- See also**
- ◆ [ls on page 296](#)
  - ◆ [unalias on page 485](#)

## array claim

Claims and names unclaimed storage volumes for a given array.

**Contexts** All contexts.

**Syntax**

```
array claim
  [-s|--storage-array] context-path
  [-m|--mapping-file] mapping file
  [-t|--tier]
  [-l|--claim]
```

**Arguments** **Required arguments**

**[-s|--storage-array] *context-path*** - \* Context path of the storage-array on which to claim storage volumes.

**Optional arguments**

**[-m|--mapping-file] *mapping file*** - Location of the name mapping file.

**[-t|--tier] *mapping file*** - Add a tier identifier to the storage volumes to be claimed.

**[-l|--claim]** - Try to claim unclaimed storage-volumes.

\* - argument is positional.

**Description** Claims and names unclaimed storage volumes for a given array.

Some storage arrays support auto-naming (EMC Symmetrix, HDS 9970/9980 and USP VM) and do not require a mapping file.

Other storage arrays require a hints file generated by the storage administrator using the array's command line. The hints file contains the device names and their World Wide Names. Refer to [Table 21, "Create hints files for storage-volume naming,"](#).

Use the **--mapping-file** argument to specify a hints file to use for naming claimed storage volumes. File names will be used to determine the array name.

Use the **--tier** argument to add a storage tier identifier in the storage-volume names.

**Example** In the following example, the array claim command:

- ◆ Claims and all storage\_volumes from the selected array
- ◆ Names the storage-volumes "Clar0400L\_<lun name>" using mapping file /tmp/Clar0400.txt
- ◆ Adds the tier identifier "L" to the storage-volume names:

```
Vplexcli: /> array claim -s FNM00083800066-0x050060160bce00f37 --mapping-file=/tmp/Clar0400.txt
--tier=L
```

**See also** ◆ [storage-volume find-array on page 464](#)

## array re-discover

Re-discovers an array, and makes the array's storage volumes visible to the VPLEX.

**Contexts** Cluster-specific context and lower.

**Syntax**  
`array re-discover`  
`[-a|--array] context-path`  
`[-c|--cluster] cluster-id`  
`[-f|--force]`

**Arguments** **Required arguments**

`[-a|--array] context-path` - \* Context path of the storage-array to re-discover.

`[-c|--cluster] cluster-id` - Cluster ID of the target cluster.

**Optional arguments**

`[-f|--force]` - Force the operation to continue without confirmation. Allows the command to be run from a non-interactive script.

\* - argument is positional.

**Description** Manually synchronizes the export state of the target device. Used in two scenarios:

- ◆ When the exported LUNs from the target array to VPLEX are modified.  
 Newer protocol-compliant SCSI devices return a notification code when the exported set changes, and may not require manual synchronization. Older devices that do not return a notification, must be manually synchronized.
- ◆ When the array is not experiencing I/O (the transport interface is idle), there is no mechanism by which to collect the notification code. In this scenario, do one of the following:
  - Wait until I/O is attempted on any of the LUNs,
  - Disruptively disconnect and reconnect the array, or
  - Use the **array rediscover** command.



### CAUTION

**This command cannot detect LUN-swapping conditions on the array(s) being re-discovered. On older configurations, this might disrupt I/O on more than the given array. If the --force argument is not used, confirmation is required.**

Use the `ll /clusters/*/storage-elements/storage-arrays/` command to display the names of storage arrays.

### Example

```
VPlexcli:/> ll /clusters/*/storage-elements/storage-arrays/
```

```
/clusters/cluster-1/storage-elements/storage-arrays:
Name                               ID                               Connectivity  Auto  Ports  Logical                               Unit
-----                               -                               -            -    -      -                               -    Count
-----                               -                               -            -    -      -                               -    -----
EMC-0x00000000192601378            0x00000000192601378  ok           -    -      0x5000097208158918, 1
   0x500009720815891d,
   0x5000097208158920,
   0x5000097208158925,
   0x5000097208158958,
```

```
VPlexcli:/> cd /clusters/cluster-1
```

```
Vplexcli:/clusters/cluster-1> array re-discover  
storage-elements/storage-arrays/EMC-0x00000000192601378  
OR:
```

```
Vplexcli:/> cd /clusters/cluster-1/storage-elements/storage-arrays/EMC-0x00000000192601378
```

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-arrays/EMC-0x00000000192601378> array  
re-discover
```

WARNING: This command cannot detect LUN-swapping conditions on the array(s) being re-discovered. LUN swapping is the swapping of LUNs on the back-end. This command cannot detect LUN swapping conditions when the number of LUNs remains the same, but the underlying actual logical units change. I/O will not be disrupted on the LUNS that do not change. Continue? (Yes/No)

**See also** ♦ [storage-volume find-array on page 464](#)

## array used-by

Displays the components that use a specified storage-array.

**Contexts** All contexts.

**Syntax** `array used-by`  
`[-a|--array] context-path`

**Arguments** `[-a|--array] context-path` - \* Specifies the storage-array for which to find users. This argument is not required if the context is the target array.

\* - argument is positional.

**Description** Displays the components (storage-volumes) that use the specified storage array.

**Example** Display the usage of components in an array from the target storage array context:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-arrays/EMC-CLARiiON-APM00050404263>
array used-by
```

```
Used-by details for storage-array EMC-CLARiiON-APM00050404263:
```

```
/clusters/cluster-1/storage-elements/extents/extent_6006016061211100363da903017ae011_1:
SV1
```

```
/clusters/cluster-1/devices/dev_clus1:
extent_SV1_1
SV1
```

```
/clusters/cluster-1/system-volumes/log1_vol:
extent_SV1_2
SV1
```

```
/clusters/cluster-1/devices/clus1_device1:
extent_SV1_3
SV1
```

```
/clusters/cluster-1/devices/clus1_dev2:
extent_SV1_4
SV1
```

```
/clusters/cluster-1/devices/device_6006016061211100d42febba1bade011_1:
extent_6006016061211100d42febba1bade011_1
VPD83T3:6006016061211100d42febba1bade011
```

```
/distributed-storage/distributed-devices/dev1_source:
dev1_source2012Feb16_191413
extent_sv1_1
sv1
```

```
/clusters/cluster-1/system-volumes/MetaVol:
VPD83T3:6006016022131300de76a5cec256df11
```

```
/clusters/cluster-1/system-volumes/MetaVol:
VPD83T3:600601606121110014da56b3b277e011
```

```
/clusters/cluster-1/system-volumes/MetaVol_backup_2012Feb13_071901:
VPD83T3:6006016061211100c4a223611bade011
```

Summary:

```
Count of storage-volumes that are not in use: 0
```

```
Count of storage-volumes that are in use: 6
```

**Example** Display the usage of components in an array from the /storage-arrays context:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-arrays> array used-by --array  
EMC-CLARiion-APM00050404263
```

```
Used-by details for storage-array EMC-CLARiion-APM00050404263:
```

```
/clusters/cluster-1/storage-elements/extents/extent_6006016061211100363da903017ae011_1:  
SV1  
.  
.  
.
```

- See also**
- ◆ [storage-volume find-array on page 464](#)
  - ◆ [storage-volume summary on page 469](#)

## authentication directory-service configure

Configures a directory service on the VPLEX cluster to authenticate users against a remote directory server.

**Contexts** All contexts.

**Syntax**

```
authentication directory-service configure
[-d|--directory-server [1|2]
[-i|--ip-address] IP address
[-b|--base-distinguished-name] "base-distinguished-name"
[-m|--map-principal] "map-principal"
[-n|--bind-distinguished-name] "bind-distinguished-name"
[-p|--bind-password]
[-s|--server-name] server-name
[-t|--connection-type] [1|2]
[-o|--port] port
[-c|--custom-attributes]
--dry-run
```

**Arguments** **Required arguments**

**[-d|--directory-server] [1|2]** - Specifies the directory server to configure on the cluster to authenticate users.

- **1** - Configures the directory service to map attributes for OpenLDAP directory with POSIX attributes.
- **2** - Configures the directory service to map attributes for Active Directory with Service For Unix (SFU 3.0) attributes.

**Note:** If option 2 (Active Directory) is selected, use the **--custom-attributes** argument to map attributes if directory server UNIX attributes are different.

**[-i|--ip-address] *IP address*** - IP address of the directory server.

**[-b|--base-distinguished-name] "*base-distinguished-name*"** - The base Distinguished Name (DN). Searches of the LDAP/AD directory start from the base DN. A DN is a sequence of relative distinguished names connected by commas. For example:

```
dc=org,dc=company,dc=com
```

The Distinguished Name must be enclosed in quotes.

**[-m|--map-principal] "*map-principal*"** - Maps a directory server Distinguished Name to VPLEX Cluster. Map-principal is a sequence of relative distinguished names connected by commas. Searches in the OpenLDAP/AD directory are restricted to the principal mapped. For example:

```
OU=eng,dc=vplex,dc=security,dc=lab,dc=emc,dc=com
```

The map-principal must be enclosed in quotes.

To include a backslash (\) in the map-principal, precede the backslash by a second backslash. For example, to specify the map-principal:

```
CN=user\ test: (IT),OU=group,DC=company,DC=com
```

Enter:

```
"CN=user\\ test: (IT),OU=group,DC=company,DC=com"
```

**[-n | --bind-distinguished-name] "bind-distinguished-name"** - The Bind Distinguished Name of the OpenLDAP/ Active Directory server. For example:

```
cn=Manager,dc=my-domain,dc=com
```

The Bind Distinguished Name must be enclosed in quotes.

**[-p | --bind-password]** - Password of Bind Distinguished Name. A prompt for the password is displayed. The password is not displayed as it is typed.

#### Optional arguments

**[-s | --server-name] server-name** - Name of the directory server. This argument is required when the **--connection-type** argument is set to 2 (SSL). If more than one server is specified, they must have the same configuration (connection type, admin account). Multiple server names must be separated by commas.

**[-t | --connection-type] {1 | 2}** - Select the cryptographic protocol to use to connect to the LDAP/ Active Directory server. Values are:

- **1** - (Default) - Use the StartTLS.
- **2** - Use the Secure Sockets Layer (SSL). If SSL is selected, use the **--server-name** argument to specify the LDAP server.

**[-o | --port] port** - Port number of the LDAP/ Active Directory server.

Range: 1- 65536.

Default: 389 when **--connection-type** is set to TLS (1).

**[-c | --custom-attributes]** - Provide custom attribute names for attribute mapping. Prompts for mapping the attribute names.

**--dry-run** - Run the command but don't do anything.

#### Description

This command configures an authentication service on the VPLEX cluster.

VPLEX supports two authentication service providers to authenticate users: OpenLDAP and Active Directory servers.

When VPLEX is configured with OpenLDAP, it uses POSIX account attribute mapping by default. When VPLEX is configured with Active Directory server, it uses SFU 3.0 attribute mapping by default.

If directory server UNIX attributes are different, use **--custom-attributes** and **--directory-server** arguments.

In order to authenticate users, a directory service must be configured on the VPLEX. Configuration includes:

- ◆ The type of directory server (OpenLDAP or Active Directory).
  - OpenLDAP by default maps POSIX attributes.
  - Active Directory by default maps SFU 3.0 attributes.
- ◆ The directory server's IP address
- ◆ Whether the TLS or SSL protocol is used
- ◆ Base Distinguished Name, for example:
 

```
dc=security,dc=orgName,dc=companyName,dc=com
```
- ◆ Bind Distinguished Name, for example:
 

```
cn=Administrator,dc=security,dc=orgName,dc=companyName,dc=com
```
- ◆ Map principal, for example:

```
ou=people,dc=security,dc=orgName,dc=companyName, dc=com
```

Best practice is to add groups rather than users. Adding groups allows multiple users to be added using one map-principal. VPLEX is abstracted from any changes (modify/delete) to the user.

**Example** Configure the Active Directory directory service on the management server:

```
VPLexcli:/> authentication directory-service configure -i 10.31.52.54 -b
"dc=vplex,dc=security,dc=sveblr,dc=lab,dc=emc,dc=com" -n
"cn=Administrator,cn=Users,dc=vplex,dc=security,dc=sveblr,dc=lab,dc=emc,dc=com" -p -d 2 -m
"cn=testUser1,cn=Users,dc=vplex,dc=security,dc=sveblr,dc=lab,dc=emc,dc=com"
```

Enter Administrator's password:

Connecting to authentication server (may take 3 minutes) ...

VPLexcli:/>

**Example** Configure the Active Directory service using custom attributes:

**Note:** Default values are displayed in brackets for each attribute. Press Enter to accept the default, or type a value and press Enter.

```
VPLexcli:/> authentication directory-service configure -i 10.31.52.189 -b
'dc=INTRANET,dc=VPLEX,dc=COM' -m 'ou=ldn' -n
'cn=Administrator,cn=Users,dc=INTRANET,dc=VPLEX,dc=COM' -d 2 -p -t 1 --server-name
svctag-4xb7mls -o 389 --custom-attributes
```

Enter Administrator's password:

Set value for posixAccount attribute [User]:

Set value for posixGroup attribute [Group]:

Set value for uid attribute [msSFU30Name]: **samaccountname**

Set value for uidNumber attribute [msSFU30UidNumber]: **uidnumber**

Set value for gidNumber attribute [msSFU30GidNumber]: **gidnumber**

Set value for loginShell attribute [msSFU30LoginShell]: **loginshell**

Set value for homeDirectory attribute [msSFU30HomeDirectory]: **unixhomedirectory**

Connecting to authentication server (may take 3 minutes) ...

VPLexcli:/>

**Example** Configure the Open LDAP service using custom attributes:

**Note:** Values for Open LDAP are case-insensitive.

```
VPLexcli:/> authentication directory-service configure -d 1 -i 10.31.52.53 -b
"dc=security,dc=sve,dc=emc,dc=com" -n "cn=Administrator,dc=security,dc=sve,dc=emc,dc=com" -m
"ou=people" -p --custom-attributes
```

Enter Administrator's password:

Set value for posixAccount attribute [posixAccount]: **posixaccount**

Set value for posixGroup attribute [posixGroup]: **posixgroup**

Set value for uid attribute [uid]: **uid**

```
Set value for uidNumber attribute [uidNumber]: uidNumber  
Set value for gidNumber attribute [gidNumber]: gidNumber  
Set value for loginShell attribute [loginShell]:  
Set value for homeDirectory attribute [homeDirectory]: homedirectory  
Connecting to authentication server (may take 3 minutes) ...  
Vplexcli:/>
```

- See also**
- ◆ [authentication directory-service map on page 38](#)
  - ◆ [authentication directory-service show on page 39](#)
  - ◆ [authentication directory-service unconfigure on page 40](#)
  - ◆ [configuration configure-auth-service on page 110](#)

## authentication directory-service map

Maps a directory server user or a user group to a VPLEX cluster.

**Contexts** All contexts.

**Syntax** `authentication directory-service map  
[-m|--map-principal] "map-principal"  
--dry-run`

**Arguments** **Required arguments**

`[-m|--map-principal] "map-principal"` - Map a directory Server user or user group to the cluster. A map-principal is a sequence of relative distinguished names connected by commas. For example:

```
OU=eng,dc=vplex,dc=security,dc=lab,dc=emc,dc=com
```

The map-principal must be enclosed in quotes.

To include a backslash (\) in the map-principal, precede the backslash by a second backslash. For example, to specify the map-principal:

```
CN=user\ test: (IT),OU=group,DC=company,DC=com
```

Enter:

```
"CN=user\\ test: (IT),OU=group,DC=company,DC=com"
```

**Optional arguments**

`--dry-run` - Run the command but don't do anything.

**Description** A directory server user is an account that can log in to the VPLEX. Each user is a member of at least one group.

Best practice is to add groups rather than users. Adding groups allows multiple users to be added using one map-principal. VPLEX is abstracted from any changes (modify/delete) to the user.

The ldap.conf file cannot exceed 4096 bytes. There is enough space in the file to map 8 - 10 principals. This command will fail with the following error if adding a principal causes the file to exceed 4096 bytes:

```
cause: Cannot map principal, limit exceeded
```

**Example** Map a directory server user group to VPLEX cluster:

```
VPllexcli:/> authentication directory-service map --map-principal "ou=group,dc=company,dc=com"
```

**See also**

- ◆ [authentication directory-service configure on page 34](#)
- ◆ [authentication directory-service show on page 39](#)
- ◆ [authentication directory-service unmap on page 41](#)

## authentication directory-service show

Displays configuration of the directory service used to authenticate users.

**Contexts** All contexts.

**Syntax** authentication directory-service show

**Example** Display the authentication service configuration:

```
Vplexcli: /> authentication directory-service show
default-authentication-service: Native VPLEX

external-authentication-service: OpenLDAP

ip: 10.31.50.102

base-dn: dc=example,dc=com

connection-type: TLS

mapped-principal: ['ou=People,dc=example,dc=com', 'uid=user1,ou=People,dc=example,dc=com']
```

**Table 2 authentication directory-service show field descriptions**

| Field                           | Description                                                                                                                                                                         |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| default-authentication-service  | Always Native VPLEX                                                                                                                                                                 |
| external-authentication-service | Directory service configured on the cluster.<br><b>OpenLDAP</b> - Lightweight Directory Access Protocol.<br><b>AD</b> - Windows Active Directory 2003 with SFU 3.5                  |
| ip                              | IP address of the directory server.                                                                                                                                                 |
| base-dn                         | Base Distinguished Name of the directory service. Components:<br><b>dc</b> - domainComponent<br><b>cn</b> - commonName<br><b>ou</b> - organizationalUnitName<br><b>uid</b> - userid |
| connection-type                 | Cryptographic protocol used to connect to the LDAP server.<br><b>TLS</b> (default) - LDAPv3 Transport Layer Security<br><b>SSL</b> - Secure Sockets Layer                           |
| uri                             | Uniform resource identifier of the directory server.                                                                                                                                |
| mapped-principal                | Users and groups mapped to directory server.                                                                                                                                        |

- See also**
- ◆ [authentication directory-service configure on page 34](#)
  - ◆ [authentication directory-service map on page 38](#)

## authentication directory-service unconfigure

Unconfigures a directory service on the VPLEX cluster.

**Contexts** All contexts.

**Syntax** authentication directory-service unconfigure  
[-f|--force]

**Arguments** **Required arguments**

None.

**Optional arguments**

[-f|--force] - Force the unconfigure without asking for confirmation.

**Description** Removes the existing directory service configuration from the VPLEX management server.

**Example** Remove the existing directory service configuration:

```
VPLexcli:/> authentication directory-service unconfigure
```

This command will unconfigure the existing directory service.. Continue? (Yes/No) **Yes**

- See also**
- ◆ [authentication directory-service configure on page 34](#)
  - ◆ [authentication directory-service show on page 39](#)
  - ◆ [authentication directory-service unmap on page 41](#)
  - ◆ [configuration configure-auth-service on page 110](#)

## authentication directory-service unmap

Unmaps the specified directory server user or group from the cluster.

**Contexts** All contexts.

**Syntax** authentication directory-service unmap  
[-m|--map-principal] "*mapped-principal*"

**Arguments** **Required arguments**

**[-m|--map-principal] "*mapped-principal*"** - Mapped directory server Distinguished Name to unmap. For example:

```
ou=eng,dc=vplex,dc=security,dc=lab,dc=emc,dc=com
```

The map-principal must be enclosed in quotes.

To include a backslash (\) in the map-principal, precede the backslash by a second backslash. For example, to specify the map-principal:

```
CN=user\ test: (IT),ou=group,DC=company,DC=com
```

Enter:

```
"CN=user\\ test: (IT),OU=group,DC=company,DC=com"
```

**Note:** The "*mapped-principal*" argument is case-sensitive and varies depending on the directory service configured:

If the directory service is OpenLDAP, entries in the mapped principal argument ('ou', 'uid', 'dc) must be specified in small letters (no capitals).

If the directory service is AD, entries in the mapped principal argument ('OU', 'CN', 'DC') must be specified in capital letters.

Use the [authentication directory-service show on page 39](#) to display the current mapped principals.

**Description** There must be at least one principal mapped. If there is only one principal, it cannot be unmapped.

**Example** Unmap a directory server group:

```
Vplexcli:/> Vplexadmin authentication directory-service unmap -m "ou=group,dc=company,dc=com"
```

- See also**
- ◆ [authentication directory-service configure on page 34](#)
  - ◆ [authentication directory-service map on page 38](#)
  - ◆ [authentication directory-service show on page 39](#)

## batch-migrate cancel

Cancels an active migration and returns the source volumes to their state before the migration.

**Contexts** All contexts.

**Syntax** `batch-migrate cancel  
[-f|--file] pathname`

**Arguments** **Required arguments**

`[-f|--file] pathname` - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Description** Attempts to cancel every migration in the specified batch file. If an error is encountered, a warning is printed to the console and the command continues until every migration has been processed.

---

**Note:** In order to re-run a canceled migration plan, the batch-migrate remove command must be used to remove the records of the migration. See [batch-migrate remove on page 51](#).

---

**Example** `Vplexcli:/data-migrations/device-migrations>  
batch-migrate cancel --file migrate.txt`

**See also**

- ◆ [batch-migrate clean on page 45](#)
- ◆ [batch-migrate commit on page 47](#)
- ◆ [batch-migrate create-plan on page 48](#)
- ◆ [batch-migrate pause on page 50](#)
- ◆ [batch-migrate remove on page 51](#)
- ◆ [batch-migrate resume on page 52](#)
- ◆ [batch-migrate start on page 53](#)
- ◆ [batch-migrate summary on page 55](#)

## batch-migrate check-plan

Checks a batch migration plan.

**Contexts** All contexts.

**Syntax** `batch-migrate check-plan [-f|--file] pathname`

**Arguments** **Required arguments**

`[-f|--file] pathname` - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is `/var/log/VPLEX/cli` on the management server.

**Description** Checks the following conditions:

- ◆ Block-size of source and target extents is equal (4 K bytes)
- ◆ Capacity of target extent is equal to, or larger than the source extent's capacity
- ◆ Device migrations:
  - Target device has no volumes on it
  - Source device has volumes on it
- ◆ Extent migrations:
  - Target extent is claimed and ready for use
  - Source extent is in use

All migration plans should be checked before being executed.



### CAUTION

**Migration of volumes in asynchronous consistency groups is not supported on volumes that are in active use. Schedule this activity as a maintenance activity to avoid DU.**

**Example** In the following example, a migration plan fails the check.

```
Vplexcli:> batch-migrate check-plan --file MigDev-test.txt
Checking migration plan file /var/log/VPLEX/cli/MigDev-test.txt.
```

```
Target device '/clusters/cluster-2/devices/dev1723_61C' has a volume.
Target device '/clusters/cluster-2/devices/dev1723_618' has a volume.
Plan-check failed, 2 problems.
```

**Example** In the following example, a migration plan passes the check.

```
Vplexcli:> batch-migrate check-plan --file migrate.txt
Checking migration plan file /temp/migration_plans/migrate.txt.
Plan-check passed.
```

- See also**
- ◆ [batch-migrate cancel on page 42](#)
  - ◆ [batch-migrate clean on page 45](#)
  - ◆ [batch-migrate commit on page 47](#)
  - ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate pause on page 50](#)

- ◆ [batch-migrate remove on page 51](#)
- ◆ [batch-migrate resume on page 52](#)
- ◆ [batch-migrate start on page 53](#)
- ◆ [batch-migrate summary on page 55](#)

## batch-migrate clean

Cleans the specified batch migration and deletes the source devices.

**Contexts** All contexts.

**Syntax**

```
batch-migrate clean
  [-f|--file] pathname
  [-e|--rename-targets]
```

**Arguments** **Required arguments**

**[-f|--file] *pathname*** - \*Directory and filename of migration plan file. relative paths can be used. if no directory is specified, the default directory is /var/log/VPlex/cli on the management server.

**Optional arguments**

**[-e|--rename-targets]** - rename the target devices and virtual volumes to the source device names.

\* argument is positional.

**Description** Dismantles the source device down to its storage volumes and unclaims the storage volumes.

- ◆ For device migrations, cleaning dismantles the source device down to its storage volumes. The storage volumes no longer in use are unclaimed.

For device migrations only, use the optional **--rename-targets** argument to rename the target device after the source device. If the target device is renamed, the virtual volume on top of it is also renamed if the virtual volume has a system-assigned default name.

Without renaming, the target devices retain their target names, which can make the relationship between volumes and devices less evident.

- ◆ For extent migrations, cleaning destroys the source extent and unclaims the underlying storage-volume if there are no extents on it.



### CAUTION

**This command must be run before the batch-migration has been removed. The command will not clean migrations that have no record in the CLI context tree.**

**Example** In the following example, source devices are torn down to their storage volumes and the target devices and volumes are renamed after the source device names:

```
Vplexcli:/> batch-migrate clean --rename-targets --file migrate.txt
```

Using migration plan file /temp/migration\_plans/migrate.txt for cleanup phase.

```
0: Deleted source extent
  /clusters/cluster-1/devices/R20061115_Symm2264_010, unclaimed its
  disks Symm2264_010
```

```
1: Deleted source extent
  /clusters/cluster-1/extents/R20061115_Symm2264_011, unclaimed its
  disks Symm2264_011
```

```
.
.
.
```

- See also**
- ◆ [batch-migrate cancel on page 42](#)
  - ◆ [batch-migrate check-plan on page 43](#)
  - ◆ [batch-migrate commit on page 47](#)
  - ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate pause on page 50](#)
  - ◆ [batch-migrate remove on page 51](#)

## batch-migrate commit

Commits the specified batch migration.

**Contexts** All contexts.

**Syntax** `batch-migrate commit`  
`[-f|--file] pathname`

**Arguments** **Required argument**  
`[-f|--file] pathname` - \*Directory and filename of migration plan file. relative paths can be used. if no directory is specified, the default directory is /var/log/VPLex/cli on the management server.

**Description** Attempts to commit every migration in the batch. Migrations in the batch cannot be committed until all the migrations are complete.

If an error is encountered, a warning is displayed and the command continues until every migration has been processed.

The batch migration process inserts a temporary RAID 1 structure above the source devices/extents with the target devices/extents as an out-of-date leg of the RAID. Migration can be understood as the synchronization of the out-of-date leg (the target).

After the migration is complete, the commit step detaches the source leg of the temporary RAID and removes the RAID.

The virtual volume, device, or extent is identical to the one before the migration except that the source device/extent is replaced with the target device/extent.

A migration must be committed in order to be cleaned.

**Note:** Use the “[batch-migrate summary](#)” command to verify that the migration has completed with no errors before committing the migration.

**Example** `Vplexcli:/> batch-migrate commit --file BSO_19`

**See also**

- ◆ [batch-migrate cancel on page 42](#)
- ◆ [batch-migrate check-plan on page 43](#)
- ◆ [batch-migrate clean on page 45](#)
- ◆ [batch-migrate create-plan on page 48](#)
- ◆ [batch-migrate remove on page 51](#)

## batch-migrate create-plan

Creates a batch migration plan file.

**Contexts** All contexts.

**Syntax**

```
batch-migrate create-plan
  [-f|--sources] local-devices
  [-t|--targets] local-devices
  [--file] pathname
  [--force]
```

**Arguments** **Required arguments**

**[-f|--sources] *local-devices*** - \* List of local-devices to migrate virtual volumes from. May contain wildcards.

**[-t|--targets] *local-devices*** - \* List of local-devices to migrate the source virtual volumes to. May contain wildcards.

**--file *pathname*** - \* Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Optional arguments**

**--force** - Forces an existing plan file with the same name to be overwritten.

\* - argument is positional.

**Description**

- ◆ The source and target extents must be typed as a comma-separated list, where each element is allowed to contain wildcards.
- ◆ If this is an *extent* migration, the source and target cluster must be the same.
- ◆ If this is a *device* migration, the source and target clusters can be different.
- ◆ The source and target can be either local-devices or extents. Mixed migrations from local-device to extent and vice versa are not allowed.
- ◆ The command attempts to create a valid migration plan from the source devices/extents to the target devices/extents.

If there are source devices/extents that cannot be included in the plan, a warning is printed to the console, but the plan is still created.

- ◆ Review the plan and make any necessary changes before starting the batch migration.

### Example: perform a batch migration

1. Create a migration plan.

Use the “[batch-migrate create-plan](#)” command to create a plan to migrate the volumes on all the devices at cluster-1 to the storage at cluster-2:

```
Vplexcli: /> batch-migrate create-plan migrate.txt --sources
/clusters/cluster-1/devices/* --targets
/clusters/cluster-2/devices/*
```

2. Use the “[batch-migrate check-plan](#)” command to check the plan:

```
Vplexcli: /> batch-migrate check-plan migrate.txt
```

If problems are found, correct the errors and re-run the command until the plan-check passes.

- Use the “[batch-migrate start](#)” command to start the migration:

```
Vplexcli:/> batch-migrate start migrate.txt
```

- Wait for the migration to finish:

Use the “[batch-migrate summary](#)” command to monitor the status of the migration:

```
Vplexcli:/> batch-migrate summary migrate.txt
```

```
Processed 10 migrations from batch migration BR0:
committed: 0
complete: 10
in-progress: 0
paused: 0
error: 0
cancelled: 0
no-record: 0
```

- When all the migrations are complete, use the “[batch-migrate commit](#)” command to commit the migration:

```
Vplexcli:/> batch-migrate commit migrate.txt
```

The source volumes will now reside on the target devices.

- Use “[batch-migrate clean](#)” command to clean the migration:

```
Vplexcli:/> batch-migrate clean --rename-targets --file migrate.txt
```

This will dismantle the source devices down to their storage volumes and rename the target devices and volumes after the source device names.

- Use the “[batch-migrate remove](#)” command to remove the record of the migration:

```
Vplexcli:/> batch-migrate remove migrate.txt
```

### Example: Pause/resume batch migration

Pause and resume an in-progress batch migration:

```
Vplexcli:/> batch-migrate pause migrate.txt
Vplexcli:/> batch-migrate resume migrate.txt
```

A batch-migration can be canceled at any-time up until the point it is committed.

Cancel and restart a batch migration:

```
Vplexcli:/> batch-migrate cancel migrate.txt
Vplexcli:/> batch-migrate remove migrate.txt
Vplexcli:/> batch-migrate start migrate.txt
```

#### See also

- ◆ [batch-migrate cancel on page 42](#)
- ◆ [batch-migrate check-plan on page 43](#)
- ◆ [batch-migrate clean on page 45](#)
- ◆ [batch-migrate commit on page 47](#)
- ◆ [batch-migrate pause on page 50](#)
- ◆ [batch-migrate remove on page 51](#)
- ◆ [batch-migrate resume on page 52](#)
- ◆ [batch-migrate start on page 53](#)
- ◆ [batch-migrate summary on page 55](#)

## batch-migrate pause

Pauses the specified batch migration.

**Contexts** All contexts.

**Syntax** `batch-migrate pause`  
`[--file] pathname`

**Arguments** **Required arguments**

**--file *pathname*** - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Description** Pauses every migration in the batch. If an error is encountered, a warning is displayed and the command continues until every migration has been processed.

Active migrations (a migration that has been started) can be paused and resumed at a later time.

Pause an active migration to release bandwidth for host I/O during periods of peak traffic.

Resume the migration during periods of low I/O.

Use the **batch-migrate pause --file *pathname*** command to pause the specified active migration.

Use the **batch-migrate resume --file *pathname*** command to resume the specified paused migration.

**Example** `Vplexcli:/> batch-migrate pause --file BS0_19`

**See also**

- ◆ [batch-migrate cancel on page 42](#)
- ◆ [batch-migrate check-plan on page 43](#)
- ◆ [batch-migrate clean on page 45](#)
- ◆ [batch-migrate commit on page 47](#)
- ◆ [batch-migrate create-plan on page 48](#)
- ◆ [batch-migrate remove on page 51](#)
- ◆ [batch-migrate resume on page 52](#)
- ◆ [batch-migrate start on page 53](#)
- ◆ [batch-migrate summary on page 55](#)

## batch-migrate remove

Removes the record of the completed batch migration.

**Contexts** All contexts.

**Syntax** `batch-migrate remove  
[--file] pathname`

**Arguments** **Required arguments**

**--file *pathname*** - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Description** Remove the migration record only if the migration has been committed or canceled. Migration records are in the /data-migrations/device-migrations context.

**Example** `Vplexcli:/data-migrations/device-migrations> batch-migrate remove  
--file migrate.txt`

or:

`Vplexcli:>batch-migrate remove /data-migrations/device-migrations  
--file migrate.txt.`

- See also**
- ◆ [batch-migrate cancel on page 42](#)
  - ◆ [batch-migrate check-plan on page 43](#)
  - ◆ [batch-migrate clean on page 45](#)
  - ◆ [batch-migrate commit on page 47](#)
  - ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate pause on page 50](#)
  - ◆ [batch-migrate resume on page 52](#)
  - ◆ [batch-migrate start on page 53](#)
  - ◆ [batch-migrate summary on page 55](#)

## batch-migrate resume

Attempts to resume every migration in the specified batch.

**Contexts** All contexts.

**Syntax** `batch-migrate resume`  
`[--file] pathname`

**Arguments** **Required arguments**

`--file pathname` - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Description** Resumes the given batch migration.

If an error is encountered, a warning is printed to the console and the command continues until every migration has been processed.

**Example** `Vplexcli:/> batch-migrate resume --file BS0_19`

**See also**

- ◆ [batch-migrate cancel on page 42](#)
- ◆ [batch-migrate check-plan on page 43](#)
- ◆ [batch-migrate clean on page 45](#)
- ◆ [batch-migrate commit on page 47](#)
- ◆ [batch-migrate create-plan on page 48](#)
- ◆ [batch-migrate pause on page 50](#)
- ◆ [batch-migrate remove on page 51](#)
- ◆ [batch-migrate start on page 53](#)
- ◆ [batch-migrate summary on page 55](#)

## batch-migrate start

Starts the specified batch migration.

**Contexts** All contexts.

**Syntax**

```
batch-migrate start
  [--file] pathname
  [-s|transfer-size] 40K - 128M
  --force
  --paused
```

**Arguments** **Required arguments**

**--file *pathname*** - \* Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Optional arguments**

**[-s | transfer-size] *size*** - Maximum number of bytes to transfer as one operation per device. Specifies the size of read sector designated for transfer in cache. Setting transfer size to a lower value implies more host I/O outside the transfer boundaries. Setting transfer size to a higher value may result in faster transfers. See [About transfer-size on page 53](#). Valid values must be a multiple of 4 K.

Range: 40 K - 128 M.

Default: 128 K.

**--force** - Do not ask for confirmation when starting individual migrations. Allows this command to be run using a non-interactive script, or to start cross-cluster device migrations in a Geo system without prompting for every device.

**--paused** - Starts the migration in a paused state.

\* - argument is positional.

**Description** Starts a migration for every source/target pair in the given migration-plan.



### **CAUTION**

**Migration of volumes in asynchronous consistency groups is not supported on volumes that are in use. Schedule this activity as a maintenance activity to avoid Data Unavailability.**

**Consider scheduling this activity during maintenance windows of low workload to reduce impact on applications and possibility of a disruption.**

If a migration fails to start, a warning is printed to the console. The command will continue until every migration item has been processed.

Individual migrations may ask for confirmation when they start. Use the **--force** argument to suppress these requests for confirmation.

**About transfer-size**

Transfer-size is the size of the region in cache used to service the migration. The area is globally locked, read at the source, and written at the target.

Transfer-size can be as small 40 K, as large as 128 M, and must be a multiple of 4 K. The default value is 128 K.

A larger transfer-size results in higher performance for the migration, but may negatively impact front-end I/O. This is especially true for VPLEX Metro migrations.

A smaller transfer-size results in lower performance for the migration, but creates less impact on front-end I/O and response times for hosts.

Set a large transfer-size for migrations when the priority is data protection or migration performance.

Set a smaller transfer-size for migrations when the priority is front-end storage response time.

Factors to consider when specifying the transfer-size:

- ◆ For VPLEX Metro configurations with narrow inter-cluster bandwidth, set the transfer size lower so the migration does not impact inter-cluster I/O.
- ◆ The region specified by transfer-size is locked during migration. Host I/O to or from that region is held. Set a smaller transfer-size during periods of high host I/O.
- ◆ When a region of data is transferred, a broadcast is sent to the system. Smaller transfer-size mean more broadcasts, slowing the migration.

**Example** `VPllexcli:/> batch-migrate start --file BS0_19 --transfer-size 1M`

- See also**
- ◆ [batch-migrate cancel on page 42](#)
  - ◆ [batch-migrate check-plan on page 43](#)
  - ◆ [batch-migrate clean on page 45](#)
  - ◆ [batch-migrate commit on page 47](#)
  - ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate pause on page 50](#)
  - ◆ [batch-migrate remove on page 51](#)
  - ◆ [batch-migrate resume on page 52](#)
  - ◆ [batch-migrate summary on page 55](#)
  - ◆ [dm migration start on page 213](#)

## batch-migrate summary

Displays a summary of the batch migration.

**Contexts** All contexts.

**Syntax** `batch-migrate summary`  
 `[--file] pathname`  
 `[-v|--verbose]`

**Arguments** **Required arguments**

**--file *pathname*** - Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is /var/log/VPLEX/cli on the management server.

**Optional arguments**

**[-v | verbose]** - In addition to the specified migration, displays a summary for any in-progress and paused migrations.

**Description** Displays a summary of the batch migration.

If the --verbose option is used, displays in the batch that are in an error state.

**Example** Display a batch migration:

```
Vplexcli: /> batch-migrate summary migrate.txt

Processed 10 migrations from batch migration migrate.txt:
committed: 0
complete: 10
in-progress: 0
paused: 0
error: 0
cancelled: 0
no-record: 0
```

**Example** Display a batch migration using the --verbose option:

**batch-migrate summary batch-migrate2.txt --verbose**

```
Command output:
source device      source      target device      target      migration      status      percentage      eta
-----
----- cluster ----- cluster      name ----- done -----
-----
templ_r1_0_cluster-1 cluster-1 templ_r1_0_cluster-2 cluster-2 BR1_0 complete 100 -
templ_r1_1_cluster-1 cluster-1 templ_r1_1_cluster-2 cluster-2 BR1_1 complete 100 -
templ_r1_2_cluster-1 cluster-1 templ_r1_2_cluster-2 cluster-2 BR1_2 complete 100 -
```

```
Processed 3 migrations from batch migration BR1:
committed: 0
complete: 3
in-progress: 0
queued: 0
paused: 0
error: 0
cancelled: 0
no-record: 0
```

**Table 3** batch migration summary field descriptions (1 of 2)

| Field                        | Description                                      |
|------------------------------|--------------------------------------------------|
| <b>--verbose output only</b> |                                                  |
| source device                | Local-device from which to migrate.              |
| source cluster               | Cluster on which source local-device is located. |

Table 3 batch migration summary field descriptions (2 of 2)

| Field                                     | Description                                                                                                        |
|-------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| target device                             | Local-device to which to migrate.                                                                                  |
| target cluster                            | Cluster on which target local-device is located.                                                                   |
| migration name                            | Names of migration files in the batch migration.                                                                   |
| status                                    | Status of the individual migration. See below for possible values.                                                 |
| eta                                       | For migrations currently being processed, the estimated time to completion.                                        |
| <b>--verbose and non --verbose output</b> |                                                                                                                    |
| Processed <i>n</i> migrations...          | Of the number of source-target pairs specified in the batch migration plan, the number that have been processed.   |
| committed                                 | Of the number of source-target pairs that have been processed, the number that have been committed.                |
| completed                                 | Of the number of source-target pairs that have been processed, the number that are complete.                       |
| in-progress                               | Of the number of source-target pairs that have been processed, the number that are in progress.                    |
| paused                                    | Of the number of source-target pairs that have been processed, the number that are paused.                         |
| error                                     | Jobs that encountered errors during processing.                                                                    |
| cancelled                                 | Of the number of source-target pairs that have been processed, the number that have been cancelled.                |
| no-record                                 | Of the number of source-target pairs that have been processed, the number that have no record in the context tree. |

**Note:** If more than 25 migrations are active at the same time, they are queued, their status is displayed as *in-progress*, and percentage-complete is displayed as ?.

- See also**
- ◆ [batch-migrate cancel on page 42](#)
  - ◆ [batch-migrate check-plan on page 43](#)
  - ◆ [batch-migrate clean on page 45](#)
  - ◆ [batch-migrate commit on page 47](#)
  - ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate pause on page 50](#)
  - ◆ [batch-migrate remove on page 51](#)
  - ◆ [batch-migrate resume on page 52](#)
  - ◆ [batch-migrate start on page 53](#)

## battery-conditioning disable

Disables battery conditioning on the specified backup battery unit(s).

**Contexts** All contexts.

**Syntax**

```
battery-conditioning disable
[-s|--sps-unit] context-path,context path,...
[-c|--all-at-cluster] cluster
[-t|--bbu-type] bbu-type
[-f|--force]
```

**Arguments** **[-s|--sps-unit] context path,context path...** \* Standby power supplies (SPS) units on which to disable battery conditioning. If this argument is used:

- The **--all-at-cluster** argument must not be specified
- The **--bbu-type** argument is ignored.

**[-c|--all-at-cluster] cluster** - \* The cluster on which to disable battery conditioning on all backup battery units. If this argument is used:

- The **--bbu-type** argument is required
- The **--sps-units** argument must not be specified.

**[-t|--bbu-type] bbu-type** - Type of battery unit on which to disable conditioning. For the current release, only standby-power-supply (sps) units are supported. This argument is ignored if the **--sps-unit** argument is specified.

**[-f|--force]** - Skips the user confirmation that appears when the battery unit on which conditioning is being disabled is currently undergoing conditioning. Allows the command to be executed from an non-interactive script.

\* - argument is positional.

**Description** Automatic battery conditioning of every SPS is enabled by default. Use this command to disable battery conditioning for all SPS units in a cluster, or a specified SPS unit.

Disabling conditioning on a unit that has a cycle in progress causes that cycle to abort, and user confirmation is required. Use the **--force** argument to skip the user confirmation.

Automatic battery conditioning must be disabled during:

- ◆ Scheduled maintenance
- ◆ System upgrades
- ◆ Unexpected and expected power outages



### CAUTION

**For all procedures that require fully operational SPS, ensure that SPS conditioning is disabled at least 6 hours in advance of the procedure. This prevents the SPS from undergoing battery conditioning during the procedure.**

**If the procedure starts when no SPS conditioning is scheduled, it is acceptable to disable SPS conditioning less than six hours in advance.**

Use the **battery-conditioning set-schedule** command to select the day the automatic monthly conditioning cycle starts.

Use the **battery-conditioning manual-cycle request** command to run an additional conditioning cycle on one or more backup battery units.

**Example** Disable battery conditioning for all SPS units on cluster-1:

```
VPllexcli:/> battery-conditioning disable --all-at-cluster /clusters/cluster-1 -t sps
Battery conditioning disabled on backup battery units 'stand-by-power-supply-b,
stand-by-power-supply-a'.
```

**Example** Disable battery conditioning for a specified SPS unit and display the change:

```
VPllexcli:/> battery-conditioning disable -s
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a
Battery conditioning disabled on backup battery units 'stand-by-power-supply-a'.

VPllexcli:/> ll
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning/
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:
Name                               Value
-----
enabled                             false
in-progress                          false
manual-cycle-requested               false
next-cycle                           Mon Dec 05 17:00:00 MST 2011
previous-cycle                       Fri Nov 25 13:25:00 MST 2011
previous-cycle-result                 PASS
```

**Example** Disable battery conditioning on a specified SPS that is currently undergoing battery conditioning:

```
VPllexcli:/> battery-conditioning disable -s
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a
The backup battery unit 'stand-by-power-supply-a' is currently undergoing a conditioning cycle.
Disabling conditioning will abort the cycle and cannot be undone.
Do you wish to disable conditioning on this unit anyway? (Yes/No) y
Battery conditioning disabled on backup battery units 'stand-by-power-supply-a'.
```

**See also**

- ◆ [battery-conditioning enable on page 59](#)
- ◆ [battery-conditioning manual-cycle cancel-request on page 61](#)
- ◆ [battery-conditioning manual-cycle request on page 63](#)
- ◆ [battery-conditioning set-schedule on page 65](#)
- ◆ [battery-conditioning summary on page 67](#)

## battery-conditioning enable

Enables conditioning on the specified backup battery unit(s).

**Contexts** All contexts.

**Syntax** `battery-conditioning enable`  
`[-s|--sps-unit] context-path,context path,...`  
`[-c|--all-at-cluster] cluster`  
`[-t|--bbu-type] bbu-type`

**Arguments** `[-s|--sps-unit] context path,context path...- *` Standby power supplies (SPS) units on which to enable battery conditioning. If this argument is used:

- The **--all-at-cluster** argument must not be specified
- The **--bbu-type** argument is ignored.

`[-c|--all-at-cluster] cluster - *` The cluster on which to enable battery conditioning on all backup battery units. If this argument is used:

- The **--bbu-type** argument is required
- The **--sps-units** argument must not be specified.

`[-t|--bbu-type] bbu-type -` Type of battery unit on which to enable conditioning. For the current release, only standby-power-supply (sps) units are supported. This argument is ignored if the **--sps-unit** argument is specified.

\* - argument is positional.

**Description** Use this command to enable battery conditioning for all standby power supply (SPS) units in a cluster, or for a specific SPS unit.

SPS battery conditioning assures that the battery in an engine's standby power supply can provide the power required to support a cache vault. A conditioning cycle consists of a 5 minute period of on-battery operation and a 6 hour period for the battery to recharge. Automatic conditioning runs every 4 weeks, one standby power supply at a time.

Automatic battery conditioning of every SPS is enabled by default.

Automatic battery conditioning must be disabled during:

- ◆ Scheduled maintenance
- ◆ System upgrades
- ◆ Unexpected and expected power outages

Use this command to re-enable battery conditioning after activities that require battery conditioning to be disabled are completed.

**Example** Enable battery conditioning on a specified SPS and display the change:

```
Vplexcli:/> battery-conditioning enable -s  

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a  

Battery conditioning enabled on backup battery units 'stand-by-power-supply-a'.
```

```
Vplexcli:/> ll  

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning/  

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:  

Name Value  

-----  

enabled true  

in-progress false
```

```

manual-cycle-requested  false
next-cycle              Mon Dec 05 17:00:00 MST 2011
previous-cycle         Fri Nov 25 13:25:00 MST 2011
previous-cycle-result  PASS

```

**Example** Enable battery conditioning on all SPS units in cluster-1 and display the change:

```

VPLEXcli:/> battery-conditioning enable --all-at-cluster cluster-1 -t sps
Battery conditioning enabled on backup battery units 'stand-by-power-supply-a,
stand-by-power-supply-b'.
VPLEXcli:/> ll /engines/*/stand-by-power-supplies/*/conditioning

```

```

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:

```

| Name                   | Value                        |
|------------------------|------------------------------|
| enabled                | true                         |
| in-progress            | false                        |
| manual-cycle-requested | false                        |
| next-cycle             | Tue Jan 03 17:00:00 MST 2012 |
| previous-cycle         | Fri Dec 16 21:31:00 MST 2011 |
| previous-cycle-result  | PASS                         |

```

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-b/conditioning:

```

| Name                   | Value                        |
|------------------------|------------------------------|
| enabled                | true                         |
| in-progress            | false                        |
| manual-cycle-requested | false                        |
| next-cycle             | Wed Jan 04 05:00:00 MST 2012 |
| previous-cycle         | Sat Dec 17 15:37:20 MST 2011 |
| previous-cycle-result  | PASS                         |

- See also**
- ◆ [battery-conditioning disable on page 57](#)
  - ◆ [battery-conditioning manual-cycle cancel-request on page 61](#)
  - ◆ [battery-conditioning manual-cycle request on page 63](#)
  - ◆ [battery-conditioning set-schedule on page 65](#)
  - ◆ [battery-conditioning summary on page 67](#)

## battery-conditioning manual-cycle cancel-request

Cancels a manually requested battery conditioning cycle on the specified backup battery unit.

**Contexts** All contexts.

**Syntax** `battery-conditioning manual-cycle cancel-request [-s|--sps-unit] context path`

**Arguments** **Required arguments**

`[-s|--sps-unit]` - Standby power supply (SPS) unit on which to cancel a previously requested battery conditioning. The full context path is required when this command is run from the engines context or higher.

**Description** Cancels a manually requested conditioning cycle on the specified backup battery unit.

Automatic battery conditioning cycles run on every SPS every 4 weeks.

Manually requested battery conditioning cycles are in addition to the automatic cycles.

Use this command to cancel a manual battery conditioning cycle.



### **IMPORTANT**

**This command does not abort a battery conditioning cycle that is underway. It cancels only a request for a manual cycle.**

**Example** Cancel a manually-scheduled SPS battery conditioning from the root context, and display the change:

```
VPllexcli:/> battery-conditioning manual-cycle cancel-request -s
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a
The manual conditioning cycle on 'stand-by-power-supply-a' has been canceled. The next
conditioning cycle will be performed on Sat Dec 31 17:00:00 MST 2011
```

```
VPllexcli:/> ll
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:
Name                               Value
-----
enabled                             true
in-progress                          false
manual-cycle-requested              false
next-cycle                          Sat Dec 31 17:00:00 MST 2011
previous-cycle                       Fri Dec 16 21:31:00 MST 2011
previous-cycle-result                PASS
```

**Example** Cancel a manually-scheduled SPS battery conditioning from the engine context:

```
VPllexcli:/engines/engine-1-1> battery-conditioning manual-cycle cancel-request -s
stand-by-power-supply-a
The manual conditioning cycle on 'stand-by-power-supply-a' has been canceled. The next
conditioning cycle will be performed on Sat Dec 31 17:00:00 MST 2011
```

**See also**

- ◆ [battery-conditioning disable on page 57](#)
- ◆ [battery-conditioning enable on page 59](#)

- ◆ [battery-conditioning manual-cycle request on page 63](#)
- ◆ [battery-conditioning set-schedule on page 65](#)
- ◆ [battery-conditioning summary on page 67](#)

## battery-conditioning manual-cycle request

Manually request a battery conditioning cycle on the specified backup battery unit.

**Contexts** All contexts.

**Syntax** `battery-conditioning manual-cycle request`  
`[-s|--sps-unit] context path`  
`[-f|--force]`

**Arguments** **Required arguments**

`[-s|--sps-unit]` - \* Standby power supply (SPS) unit(s) on which to request the battery conditioning cycle. The full context path is required when this command is run from the engines context or higher.

**Optional arguments**

`[-f|--force]` - Forces the requested battery conditioning cycle to be scheduled without confirmation if the unit is currently in a conditioning cycle. Allows this command to be run from non-interactive scripts.

\* - argument is positional.

**Description** Requests a conditioning cycle on the specified backup battery unit, and displays the time the cycle is scheduled to start. The requested battery conditioning cycle is scheduled at the soonest available time slot for the specified unit.

If the specified unit is currently undergoing a conditioning cycle, this command requests an additional cycle to run at the next available time slot.

If battery conditioning is disabled, the manually requested cycle will not run.

Use this command to manually schedule a battery conditioning cycle when automatic conditioning has been disabled in order to perform maintenance, upgrade the system, or shut down power.

The conditioning cycle invoked by this command will run in the next 6-hour window available for the selected unit.

**Note:** Scheduling a manual conditioning cycle is while a conditioning cycle is already in progress contributes to shortened battery life and is not recommended.

**Example** Schedule a manual SPS battery conditioning cycle from the root context and display the change:

```
Vplexcli:/> battery-conditioning manual-cycle request -s  

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a  

A manual conditioning cycle will be performed on 'stand-by-power-supply-a' on Tue Dec 20  

23:00:00 MST 2011.
```

```
Vplexcli:/>  

ll /engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning  

/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:  

Name Value  

-----  

enabled true  

in-progress false  

manual-cycle-requested true  

next-cycle Tue Dec 20 23:00:00 MST 2011  

previous-cycle Fri Dec 16 21:31:00 MST 2011  

previous-cycle-result PASS
```

**Example** Schedule a manual SPS battery conditioning from the engines/*engine* context when a conditioning cycle is already underway:

```
Vplexcli:/engines/engine-1-1/> battery-conditioning manual-cycle request  
stand-by-power-supply-a/
```

The backup battery unit 'stand-by-power-supply-a' is currently undergoing a conditioning cycle. Scheduling a manual cycle now is unnecessary and discouraged as it will contribute to over-conditioning and a shortened battery life.

Do you wish to manually schedule a conditioning cycle on this unit anyway? (Yes/No) **y**  
A manual conditioning cycle will be performed on 'stand-by-power-supply-a' on Fri Nov 25 13:25:00 MST 2011.

- See also**
- ◆ [battery-conditioning disable on page 57](#)
  - ◆ [battery-conditioning enable on page 59](#)
  - ◆ [battery-conditioning manual-cycle cancel-request on page 61](#)
  - ◆ [battery-conditioning set-schedule on page 65](#)
  - ◆ [battery-conditioning summary on page 67](#)

## battery-conditioning set-schedule

Set the battery conditioning schedule (day of week) for backup battery units on a cluster.

**Contexts** All contexts.

**Syntax** `battery-conditioning set-schedule`  
`[-t|--bbu-type] bbu-type`  
`[-d|--day-of-week] [sunday|monday|...|saturday]`  
`[-c|--cluster] cluster`

**Arguments** **Required arguments**

`[-t|--bbu-type] bbu-type` - \* Type of battery backup unit to be conditioned.

**Note:** In the current release, the only *bbu-type* supported is `sps`.

`[-d|--day-of-week]` - \* Day of the week on which to run the battery conditioning. Valid values are: `sunday`, `monday`, `tuesday`, `wednesday`, etc.

`[-c|--cluster]` - \* Cluster on which to set the battery conditioning schedule.

\* - argument is positional.

**Description** Sets the day of week when the battery conditioning cycle is started on all backup battery units (*bbu*) on a cluster.

The time of day the conditioning cycle runs on an individual backup battery unit is scheduled by VPLEX.

SPS battery conditioning assures that the battery in an engine's standby power supply can provide the power required to support a cache vault. A conditioning cycle consists of a 5 minute period of on-battery operation and a 6 hour period for the battery to recharge. Automatic conditioning runs every 4 weeks, one standby power supply at a time.

Automatic battery conditioning of every SPS is enabled by default.

Use this command to set the day of the week on which the battery conditioning cycle for each SPS unit (one at a time) begins.

**Example** Set the start day of the battery conditioning cycle for all SPS units in cluster-1 to Saturday and display the change:

```
Vplexcli:/> battery-conditioning set-schedule -t sps -d saturday -c cluster-1
Battery conditioning schedule for sps units on cluster 'cluster-1' successfully set to 'saturday'.
```

```
Vplexcli:/> ll /engines/*/stand-by-power-supplies/*/conditioning/
```

```
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-a/conditioning:
```

| Name                   | Value                        |
|------------------------|------------------------------|
| enabled                | true                         |
| in-progress            | false                        |
| manual-cycle-requested | false                        |
| next-cycle             | Fri Feb 03 17:00:00 MST 2012 |
| previous-cycle         | Fri Dec 16 21:31:00 MST 2011 |
| previous-cycle-result  | PASS                         |

```
/engines/engine-1-1/stand-by-power-supplies/stand-by-power-supply-b/conditioning:
```

```

Name                               Value
-----
enabled                             true
in-progress                          false
manual-cycle-requested               false
next-cycle                           Sat Feb 04 05:00:00 MST 2012
previous-cycle                       Sat Dec 17 15:37:20 MST 2011
previous-cycle-result                 PASS

```

**Table 4** battery conditioning field descriptions

| Field                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| enabled                | Whether battery conditioning is enabled on the backup battery unit.                                                                                                                                                                                                                                                                                                                                                                                              |
| in-progress            | Whether a battery conditioning cycle (either manually requested or automatic) is currently underway on the backup battery unit.                                                                                                                                                                                                                                                                                                                                  |
| manual-cycle-requested | Whether a manually requested battery conditioning cycle is scheduled on the backup battery unit.                                                                                                                                                                                                                                                                                                                                                                 |
| next-cycle             | The date and time of the next battery conditioning cycle (either manually requested or automatic) on the backup battery unit.<br><br><b>Note:</b> If the engines have a different time zone setting than the management server, and the date and time selected for the start of battery conditioning is within the difference between the two time zones, the day selected by the <b>--day-of-week</b> argument may not be the same day displayed in this field. |
| previous-cycle         | The date and time of the previous battery conditioning cycle (either manually requested or automatic) on the backup battery unit. See Note in "next-cycle".                                                                                                                                                                                                                                                                                                      |
| previous-cycle-result  | <b>ABORTED</b> - The previous battery conditioning cycle was stopped while it was underway.<br><b>PASS</b> - The previous battery conditioning cycle was successful.<br><b>SKIPPED</b> - The previous cycle was not performed because battery conditioning was disabled.<br><b>UNKNOWN</b> - The time of the previous battery conditioning cycle cannot be determined or no cycle has ever run on the backup battery unit.                                       |

- See also**
- ◆ [battery-conditioning disable on page 57](#)
  - ◆ [battery-conditioning enable on page 59](#)
  - ◆ [battery-conditioning manual-cycle cancel-request on page 61](#)
  - ◆ [battery-conditioning manual-cycle request on page 63](#)
  - ◆ [battery-conditioning summary on page 67](#)

## battery-conditioning summary

Displays a summary of the battery conditioning schedule for all devices, grouped by type and cluster

**Contexts** All contexts.

**Syntax** battery-conditioning summary

**Arguments** Required arguments

None.

**Description** Displays a summary of the conditioning schedule for all devices, grouped by type and cluster.

**Example** Display battery conditioning schedule for all devices:

```
VPllexcli:/> battery-conditioning summary
Standby Power Supply Units
```

| Cluster   | Owner      | Unit                    | Enabled | In Progress | Next Cycle                   | Previous Cycle               | Previous Result | Manual Cycle Requested | Schedule |
|-----------|------------|-------------------------|---------|-------------|------------------------------|------------------------------|-----------------|------------------------|----------|
| cluster-1 | engine-1-1 | stand-by-power-supply-a | true    | false       | Mon Mar 05 00:00:00 UTC 2012 | Tue Feb 07 08:04:09 UTC 2012 | PASS            | false                  | monday   |
|           | engine-1-1 | stand-by-power-supply-b | true    | false       | Mon Mar 05 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-1-2 | stand-by-power-supply-a | true    | false       | Mon Feb 13 00:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-1-2 | stand-by-power-supply-b | true    | false       | Mon Feb 13 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-1-3 | stand-by-power-supply-a | true    | false       | Mon Feb 20 00:00:00 UTC 2012 |                              | unknown         | false                  | friday   |
|           | engine-1-3 | stand-by-power-supply-b | true    | false       | Mon Feb 20 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-1-4 | stand-by-power-supply-a | true    | false       | Mon Feb 27 00:00:00 UTC 2012 | Mon Feb 06 12:04:08 UTC 2012 | PASS            | false                  | monday   |
|           | engine-1-4 | stand-by-power-supply-b | true    | false       | Mon Feb 27 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |

Warning: the conditioning schedule is not consistent for all units at this cluster.

| Cluster   | Owner      | Unit                    | Enabled | In Progress | Next Cycle                   | Previous Cycle               | Previous Result | Manual Cycle Requested | Schedule |
|-----------|------------|-------------------------|---------|-------------|------------------------------|------------------------------|-----------------|------------------------|----------|
| cluster-2 | engine-2-1 | stand-by-power-supply-a | true    | false       | Mon Mar 05 00:00:00 UTC 2012 | Mon Feb 06 12:04:05 UTC 2012 | PASS            | false                  | monday   |
|           | engine-2-1 | stand-by-power-supply-b | true    | false       | Mon Mar 05 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-2-2 | stand-by-power-supply-a | true    | false       | Mon Feb 13 00:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-2-2 | stand-by-power-supply-b | true    | false       | Mon Feb 13 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-2-3 | stand-by-power-supply-a | true    | false       | Mon Feb 20 00:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-2-3 | stand-by-power-supply-b | true    | false       | Mon Feb 20 12:00:00 UTC 2012 |                              | unknown         | false                  | monday   |
|           | engine-2-4 | stand-by-power-supply-a | true    | false       | Mon Feb 27 00:00:00 UTC 2012 | Mon Feb 06 12:04:12 UTC 2012 | PASS            | false                  | monday   |
|           | engine-2-4 | stand-by-power-supply-b | true    | false       | Mon Feb 27 12:00:00 UTC 2012 | Tue Feb 07 18:05:04 UTC 2012 | PASS            | false                  | monday   |

No units currently have conditioning cycles in progress.

Units stand-by-power-supply-a@engine-1-2, stand-by-power-supply-a@engine-2-2 are next to be cycled at Mon Feb 13 00:00:00 UTC 2012

**Table 5** battery conditioning summary field descriptions (1 of 2)

| Field                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enabled                | Whether battery conditioning is enabled on the backup battery unit.                                                                                                                                                                                                                                                                                                                                                                                              |
| In-Progress            | Whether a battery conditioning cycle (either manually requested or automatic) is currently underway on the backup battery unit.                                                                                                                                                                                                                                                                                                                                  |
| Manual Cycle Requested | Whether a manually requested battery conditioning cycle is scheduled on the backup battery unit.                                                                                                                                                                                                                                                                                                                                                                 |
| Next Cycle             | The date and time of the next battery conditioning cycle (either manually requested or automatic) on the backup battery unit.<br><br><b>Note:</b> If the engines have a different time zone setting than the management server, and the date and time selected for the start of battery conditioning is within the difference between the two time zones, the day selected by the <b>--day-of-week</b> argument may not be the same day displayed in this field. |
| Owner                  | VPLEX engine associated with the specified battery.                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 5 battery conditioning summary field descriptions (2 of 2)**

| Field                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Previous Cycle        | The date and time of the previous battery conditioning cycle (either manually requested or automatic) on the backup battery unit. See Note in <a href="#">“Next Cycle”</a> .                                                                                                                                                                                                                                               |
| Previous Cycle Result | <b>ABORTED</b> - The previous battery conditioning cycle was stopped while it was underway.<br><b>PASS</b> - The previous battery conditioning cycle was successful.<br><b>SKIPPED</b> - The previous cycle was not performed because battery conditioning was disabled.<br><b>UNKNOWN</b> - The time of the previous battery conditioning cycle cannot be determined or no cycle has ever run on the backup battery unit. |
| Schedule              | Day of the week on which day of week when the battery conditioning cycle is started on all backup battery units on a cluster. Configured using the <a href="#">“battery-conditioning set-schedule”</a> command.                                                                                                                                                                                                            |
| Unit                  | ID of the power supply where the specified battery is installed.                                                                                                                                                                                                                                                                                                                                                           |

- See also**
- ◆ [battery-conditioning disable on page 57](#)
  - ◆ [battery-conditioning enable on page 59](#)
  - ◆ [battery-conditioning manual-cycle cancel-request on page 61](#)
  - ◆ [battery-conditioning manual-cycle request on page 63](#)
  - ◆ [battery-conditioning set-schedule on page 65](#)

## capture begin

Begins a capture session.

**Contexts** All contexts.

**Syntax**

```
capture begin
  [-s|session] session name
  [-c|capture-directory] capture-directory]
```

**Arguments** **Required arguments**

**[-s | --session] session name** - \* Name of capture session. Output files from the capture session are named using this value.

**[-c | --capture-directory] directory** - \* Pathname for the capture directory.

Default capture directory: /var/log/VPLEX/cli/capture

\* - argument is positional.

**Description** The session captures saves all the stdin, stdout, stderr, and session I/O streams to 4 files:

- ◆ *session name-session.txt* - Output of command(s) issued during the capture session.
- ◆ *session name-stdin.txt* - CLI command(s) input during the capture session.
- ◆ *session name-stdout.txt* - Output of command(s) issued during the capture session.
- ◆ *session name-stderr.txt* - Status messages generated during the capture session.

**Note:** Raw tty escape sequences are not captured. Use the `--capture shell` option to capture the entire session including the raw tty sequences.

Capture sessions can have nested capture sessions but only the capture session at the top of the stack is active.

Use the “[capture end](#)” command to end the capture session.

Use the “[capture replay](#)” command to resubmit the captured input to the shell.

**Example** In the following example, the `capture begin` command starts a capture session named TestCapture. Because no directory is specified, output files are placed in the /var/log/VPLEX/cli/capture directory on the management server.

```
Vplexcli:/> capture begin TestCapture
# capture begin TestCapture
```

```
Vplexcli:/>
```

- See also**
- ◆ [capture end on page 70](#)
  - ◆ [capture pause on page 71](#)
  - ◆ [capture replay on page 72](#)
  - ◆ [capture resume on page 73](#)

---

## capture end

Ends the current capture session and removes it from the session capture stack.

**Contexts** All contexts.

**Syntax** `capture end`

**Description** The session at the top of the stack becomes the active capture session.

**Example**

```
VPllexcli:/clusters/cluster-1> capture end  
# capture end TestCapture
```

```
VPllexcli:/clusters/cluster-1>
```

**See also**

- ◆ [capture begin on page 69](#)
- ◆ [capture pause on page 71](#)
- ◆ [capture replay on page 72](#)
- ◆ [capture resume on page 73](#)

---

## capture pause

Pauses the current capture session.

**Contexts** All contexts.

**Syntax** `capture pause`

**Description** Pause/resume operates only on the current capture session.

**Example** `Vplexcli: /> capture pause`

**See also**

- ◆ [capture begin on page 69](#)
- ◆ [capture end on page 70](#)
- ◆ [capture replay on page 72](#)
- ◆ [capture resume on page 73](#)

## capture replay

Replays a previously captured session.

**Contexts** All contexts.

**Syntax**

```
capture replay
  [-s|-session] session name
  [-c|--capture-directory] directory
```

**Arguments** **Required arguments**

**[-s|--session] *session name*** - \* Name of existing capture session.

**[-c|--capture-directory] *directory*** - \* Directory where existing captured session is located.

Default directory: /var/log/VPLEX/cli/capture/recapture

\* - argument is positional.

**Description** Replays the commands in the stdin.txt file from the specified capture session.

Output of the replayed capture session is written to the /var/log/VPLEX/cli/capture/recapture directory on the management server.

Output is the same four files created by **capture begin**.

**Example** VPLEXcli:/> **capture replay TestCapture**

Attributes:

| Name              | Value |
|-------------------|-------|
| allow-auto-join   | true  |
| auto-expel-count  | 0     |
| auto-expel-period | 0     |
| .                 | .     |
| .                 | .     |

**See also**

- ◆ [capture begin on page 69](#)
- ◆ [capture end on page 70](#)
- ◆ [capture pause on page 71](#)
- ◆ [capture resume on page 73](#)

---

## capture resume

Resumes the current capture session.

**Contexts** All contexts.

**Syntax** `capture resume`

**Description** Pause/resume operates only on the current capture session.

**Example** `Vplexcli: /> capture resume`

**See also**

- ◆ [capture begin on page 69](#)
- ◆ [capture end on page 70](#)
- ◆ [capture pause on page 71](#)
- ◆ [capture replay on page 72](#)

## cd

Changes the working directory.

**Contexts** All contexts.

**Arguments** **Required arguments**

None.

**Optional arguments**

**cd <ENTER>** - Change to the root directory.

**cd ...** - Change to the root directory.

**cd ..** - Change to the context immediately above the current context.

**cd *context*** - Change to the specified context.

**Description** Use the **cd** command with no arguments or followed by three periods (**cd...**) to return to the root context.

Use the **cd** command followed by two periods (**cd..**) to return to the context immediately above the current context.

Use the **cd** command followed by a dash (**cd -**) to return to the previous context.

To navigate directly to a context from any other context, use the **cd** command and specify the context path.

**cd -c** - and **cd --context** - are not supported.

**Example** Return to the root context:

```
Vplexcli:/engines/engine-1-1/fans> cd
```

```
Vplexcli: />
```

Return to the context immediately above the current context:

```
Vplexcli:/monitoring/directors/director-1-1-B> cd ..
```

```
Vplexcli:/monitoring/directors>
```

Navigate directly to a context from any other context:

```
Vplexcli:/engines/engine-2-1/fans> cd /engines/engine-1-1/fans/
```

## chart create

Creates a chart based on a CSV file produced by the report command.

**Contexts** All contexts.

**Syntax**

```
chart create
  [--input] "input file"
  [--output] "output file"
  [--series] series column
  [--range] series range
  [--domain] domain column
  [--width] chart width
  [--height] chart height
  [--aggregate] aggregate-series-name
```

**Arguments** Required arguments

**Note:** All arguments are positional.

**--input** *input file* - CSV file to read data from, enclosed in quotes.

**--output** *output file* - PNG file to save chart to, enclosed in quotes.

**--series** *column* - The column in the CSV file to use as series.

**--range** *column* - The column in the CSV file to use as range.

**--domain** *column* - The column in the CSV file to use as domain.

**--width** *nn* - The width of chart graphic.

Range: 64-2048.

Default: 500.

**--height** *nn* - The height of the chart graphic.

Range: 64-2048.

Default: 500.

**--aggregate** *aggregate-series-name* - Name of aggregate series. If not specified, no aggregate series is displayed.

**Example** In the following example:

- ◆ The **exit** command exits the CLI and returns to the management server.
- ◆ The **cd** command navigates to the reports directory.
- ◆ The **tail** command displays the .csv file created using the **report capacity clusters** command.
- ◆ The **vplexcli** and login commands return to the CLI.
- ◆ The **chart create** command creates a chart.
- ◆ The **exit** command exits the CLI and returns to the management server.
- ◆ The **cd** command navigates to the reports directory.
- ◆ The **ls** command verifies that the .png file was created.

```
Vplexcli:/> exit
Connection closed by foreign host.
service@ManagementServer:~> cd /var/log/Vplex/cli/reports
```

```

service@ManagementServer:/var/log/Vplex/cli/reports> tail CapacityClusters.csv
Time, Cluster, Unclaimed disk capacity (GiB), Unclaimed storage_volumes, Claimed disk
capacity(GiB), Claimed storage_ volumes, Used storage-volume capacity
(GiB), Used storage_volumes, Unexported volume capacity (GiB), Unexported volum
es, Exported volume capacity (GiB), Exported volumes
2010-06-21 15:59:39, cluster-1, 5705.13, 341, 7947.68, 492, 360.04, 15, 3.00, 3, 2201.47, 27
.
.
.
service@ManagementServer:~> vplexcli
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.

Enter User Name: service

Password: :
creating logfile:/var/log/Vplex/cli/session.log_service_localhost_ T28921_20101020175912

Vplexcli:> chart create "CapacityClusters.csv" "CapacityClusters.png" 1 2 0 500 500
Vplexcli:/> exit
Connection closed by foreign host.
service@ManagementServer:~> cd /var/log/Vplex/cli/reports
service@ManagementServer:/var/log/Vplex/cli/reports> ll
total 48
.
.
.
-rw-r--r-- 1 service users 844 2010-07-19 15:55 CapacityClusters.csv
-rw-r--r-- 1 service users 18825 2010-07-19 15:56 CapacityClusters.png
.
.
.

```

- See also**
- ◆ [report aggregate-monitors on page 386](#)
  - ◆ [report capacity-arrays on page 387](#)
  - ◆ [report capacity-clusters on page 390](#)
  - ◆ [report capacity-hosts on page 391](#)
  - ◆ [report create-monitors on page 392](#)

## cluster add

Adds a cluster to a running VPLEX.

**Contexts** All contexts.

**Syntax**

```
cluster add
  [-c|--cluster] context path
  [-t|--to] cluster
  [-f|--force]
```

**Arguments** **Required arguments**

**[-c|--cluster] context path** - \* Cluster to add.

**[-t|--to] cluster cluster** - \* Cluster to which the given cluster will be added. This is only necessary if the system cannot be automatically determined.

**Optional arguments**

**[-f|--force]** - Forces the cluster addition to proceed even if conditions are not optimal.

\* - argument is positional.

**Description** A cluster must be added in order to communicate with other clusters in a VPLEX.

Use the **--to** argument:

- ◆ During system bring-up when no clusters have yet been told about other clusters. In this scenario, any cluster can be used as the system representative.
- ◆ Multiple systems have been detected. Connection to multiple systems, is not supported.

If there only one system actually present, but it has split into islands due to connectivity problems, it is highly advisable to repair the problems before proceeding. Add the given cluster to each island separately.

If the intention is to merge two existing systems, break up one of the systems and add it to the other system cluster-by-cluster.

**Example** In the following example:

- ◆ The **cluster add** command adds two clusters.
- ◆ The **cluster summary** command verifies that the two clusters have the same island ID:

```
VPlexcli:/clusters/> cluster add cluster-1 cluster-2
```

```
VPlexcli:/clusters> cluster summary
```

```
Clusters:
Name      Cluster ID  Connected  Expelled  Operational Status  Health State
-----
cluster-1 1          true       false     ok                 ok
cluster-2 2          true       false     ok                 ok
```

```
Islands:
Island ID  Clusters
-----
1          cluster-1, cluster-2
```

- See also**
- ◆ [cluster expel on page 83](#)
  - ◆ [cluster status on page 89](#)

- ◆ cluster summary on page 93

## cluster cacheflush

Flushes the cache on directors at the specified clusters to the back-end storage volumes.

**Contexts** All contexts.  
In /clusters context, command is **cacheflush**.

**Syntax**

```
cluster cacheflush
  [-e|--sequential]
  [-c|--clusters] cluster,cluster
  [-v|--volumes] volumes
  --verbose
```

**Arguments** **Required arguments**

**[-c|--clusters] *clusters,cluster*** - Flushes the cache for every exported virtual volume of every director at the specified clusters. Entered as wildcard patterns.

**[-v|--volumes] *volumes*** - Flushes the cache only for the specified list of virtual volumes. Entries must be separated by commas. Wildcard patterns (CLI glob patterns) are allowed.

**Optional arguments**

**[-e|--sequential]** - Flushes the cache of multiple directors sequentially. Default is to flush the caches in parallel.

**--verbose** - Displays a progress report during the flush. Default is to display no output if the run is successful.

### Description



#### **IMPORTANT**

**The CLI must be connected to a director before the cache can be flushed. Only exported virtual volumes can be flushed.**

When executed from a specific cluster context, this command flushes the cache of the directors at the current cluster.

### Example

```
Vplexcli:/clusters/> cacheflush --clusters /clusters/* --verbose --sequential
```

```
Flushing director1 director2 ...
  Flushing logging_logging1_vol @ director1... done
  Flushing regression_dR1a_12_vol @ director1... done
  Flushing regression_dR1b_12_vol @ director1... done
  .
  .
  .
  Flushing regression_r0a_1_vol @ director1... volume not exported, skipping.
  Flushing regression_r0b_1_vol @ director1... volume not exported, skipping.
  Flushing regression_r0c_1_vol @ director1... volume not exported, skipping.
  .
  .
  .
  Flushing logging_logging1_vol @ director2... done
  Flushing regression_dR1a_12_vol @ director2... done
  Flushing regression_dR1b_12_vol @ director2... done
  .
  .
  .
```

```
Flushing regression_r1c_1_vol @ director2... volume not exported, skipping.  
Flushing regression_r1d_1_vol @ director2... volume not exported, skipping.  
Flushing regression_r1e_1_vol @ director2... volume not exported, skipping.
```

- See also**
- ◆ [cluster status on page 89](#)
  - ◆ [cluster summary on page 93](#)

## cluster configdump

Dumps cluster configuration in an XML format, optionally directing it to a file.

**Contexts** All contexts.

In /clusters context, command is **configdump**.

**Syntax**

```
cluster configdump
  [-c|--cluster] cluster
  [-d|--dtdOnly]
  [-f|--file] filename
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**[-c|--clusters] *cluster*** - Dump configuration information for only the specified cluster.

**[-d|--dtdOnly]** - Print only the Document Type Definitions (DTD) document.

**[-f|--file] *filename*** - Direct the configdump output to the specified file.

Default location for the output file on the management server is: /var/log/VPLEX/cli.

**Description** Dumped data includes:

- ◆ I/O port configurations
- ◆ Disk information, including paths from the directors to the storage volumes
- ◆ Device configuration and capacity
- ◆ Volume configuration
- ◆ Initiators
- ◆ View configuration
- ◆ System-volume information

The XML output includes the DTD to validate the content.

**Example** Dump cluster-1's configuration to an .xml file:

```
Vplexcli:/clusters/>configdump -c cluster-1 -f cluster1_config.xml
```

**Example** Dump the configuration at cluster-1, navigate to the cli directory on the management server, and display the file:

```
Vplexcli:/clusters> configdump --verbose --file /var/log/VPLEX/cli/config-dump-cluster-1.txt --cluster cluster-1
Vplexcli:/clusters> exit
Connection closed by foreign host.
service@ManagementServer:~> cd /var/log/VPLEX/cli
service@ManagementServer:/var/log/VPLEX/cli>tail config-dump-cluster-1.txt
  </views>
  <system-volumes>
    <meta-volumes>
      <meta-volume active="true" block-count="23592704" block-size="4096B" geometry="raid-1"
locality="local" name="metadata_1" operational-status="ok" ready="true" rebuild-allowed="true"
size="96635715584B" system-id="metadata_1"/>
    </meta-volumes>
  </system-volumes>
  <logging-volumes>
```

```
<logging-volume block-count="20971520" block-size="4096B" geometry="raid-0"
locality="local" name="logging_1_vol" operational-status="ok" size="85899345920B"
system-id="logging_logging_1_vol"/>
</logging-volumes>
</system-volumes>
```

.  
.  
.

- See also**
- ◆ [collect-diagnostics](#) on page 106
  - ◆ [director appcon](#) on page 189
  - ◆ [getsysinfo](#) on page 272
  - ◆ [sms dump](#) on page 448

## cluster expel

Expels a cluster from its current island.

**Contexts** All contexts.  
In /clusters context, command is **expel**.

**Syntax**

```
cluster expel
  [-c|--cluster] cluster
  [-f|--force]
```

**Arguments** **Required arguments**  
[-c|--clusters] *cluster* - \* The cluster to expel.  
[-f|--force] - \* Forces the cluster to be expelled.  
\* - argument is positional.

**Description** Cluster expulsion prevents a cluster from participating in a VPLEX. Expel a cluster when:

- ◆ The cluster is experiencing undiagnosed problems.
- ◆ To prepare for scheduled outage.
- ◆ The target cluster, or the WAN over which the rest of the system communicates with it, is going to be inoperable for a while.
- ◆ An unstable inter-cluster link impacts performance.

An expelled cluster is still physically connected to the VPLEX, but not logically connected.

The **--force** argument is required for the command to complete.

Use the **cluster unexpel** command to allow the cluster to rejoin the island.

**Example** In the following example:

- ◆ The **cluster expel** command expels the cluster.
- ◆ The **cluster summary** and **cluster status** commands verify the change.

```
Vplexcli:./> cluster expel cluster-1 --force
Cluster 'cluster-1' has been successfully expelled.
```

```
Vplexcli:./> cluster summary
```

Clusters:

| Name      | Cluster ID | Connected | Expelled | Operational Status | Health State |
|-----------|------------|-----------|----------|--------------------|--------------|
| cluster-1 | 1          | true      | true     | isolated           | degraded     |
| cluster-2 | 2          | true      | false    | degraded           | degraded     |

Islands:

| Island ID | Clusters  |
|-----------|-----------|
| 1         | cluster-1 |
| 2         | cluster-2 |

```
Vplexcli:./> cluster status
```

```
Cluster cluster-1
operational-status:      isolated
```

```
transitioning-indications: suspended volumes,expelled
transitioning-progress:
health-state: degraded
health-indications: 1 suspended Devices
```

```
Cluster cluster-2
operational-status: degraded
transitioning-indications: suspended exports,suspended volumes
transitioning-progress:
health-state: degraded
health-indications: 2 suspended Devices
```

**See also** ♦ [cluster unexpel on page 95](#)

## cluster forget

Tells VPLEX and Unisphere for VPLEX to forget the specified cluster.

**Contexts** All contexts.

In /clusters context, command is **forget**.

**Syntax**

```
cluster forget
  [-c|--cluster] context path
  [-d|--disconnect]
  [-f|--force]
```

**Arguments** **Required arguments**

**[-c|--clusters] context path** - \* Cluster to forget.

**Optional arguments**

**[-d|--disconnect]** - Disconnect from all directors in the given cluster and remove the cluster from the context tree after the operation is complete.

**[-f|--force]** - Force the operation to continue without confirmation.

\* - argument is positional.

**Description** Removes all references to the specified cluster from the context tree.

The prerequisites for forgetting a cluster are as follows:

- ◆ The target cluster can not be in contact with other connected clusters.
- ◆ The Unisphere for VPLEX cannot be connected to the target cluster.
- ◆ Detach all distributed devices with legs at the target cluster (there must be no distributed devices with legs on the target cluster).
- ◆ No rule sets that affect the target cluster.
- ◆ No globally visible devices at the target cluster.

Use the following steps to forget a cluster:

1. If connected, use the **cluster forget** command on the target cluster to forget the other clusters.
2. Use the **cluster forget** command on all other clusters to forget the target cluster.

This command does not work if the clusters have lost communications with each other. If a cluster is down, destroyed, or removed, use the **cluster expel** command to expel it.

**Example**

```
Vplexcli:/clusters/> cluster forget --cluster cluster-1 --disocnnect
--force
```

**See also**

- ◆ [cluster add on page 77](#)
- ◆ [cluster expel on page 83](#)
- ◆ [cluster status on page 89](#)
- ◆ [cluster unexpel on page 95](#)

## cluster shutdown

Starts the orderly shutdown of all directors at a single cluster.

**Contexts** All contexts.

In /clusters context, command is **shutdown**.

**Syntax**

```
cluster shutdown
  [-c|--cluster] context path
  --force
```

**Arguments** **Required arguments**

**[-c|--cluster] context path** - Cluster to shut down.

**[-f|--force]** - Forces the shutdown to proceed.



### **WARNING**

*Shutting down a VPLEX cluster may cause data unavailability. Please refer to the VPLEX Procedure Generator for the recommended procedure to shut down a cluster.*

**Description** Shuts down the cluster firmware.

**Note:** Does not shut down the operating system on the cluster.

Use this command as an alternative to manually shutting down the directors in a cluster. When shutting down multiple clusters:

- ◆ Shut each cluster down one at a time.
- ◆ Verify that each cluster has completed shutdown prior to shutting down the next one.

If shutting down multiple clusters, refer to the VPLEX Procedure Generator for the recommended procedure for shutting down both clusters.

When a cluster completes shutting down, the following log message is generated for each director at the cluster:

'Director shutdown complete (cluster shutdown)'

**Example** Start the shutdown of the specified cluster:

```
VPlxcli:/> cluster shutdown -c cluster-1 --force
Status      Description
-----
Started.    Shutdown started.
```

**Example** In the following example:

- ◆ The **cluster shutdown** command without the **--force** argument starts the shutdown of the specified cluster.  
Because the **--force** argument was not used, a prompt to continue is displayed.
- ◆ The **cluster summary** commands display the transition to shutdown.
- ◆ The **ll** command in clusters/cluster-n context displays the shutdown cluster.

```
VPlxcli:/> cluster shutdown -c cluster-1
```

Warning: Shutting down a VPLEX cluster may cause data unavailability. Please refer to the VPLEX documentation for the recommended procedure for shutting down a cluster. Do you wish to proceed ? (Yes/No) **Yes**

```
Status      Description
-----
Started.    Shutdown started.
```

VPLEXcli:/> **cluster summary**

Clusters:

| Name      | Cluster ID | Connected | Expelled | Operational Status | Health State |
|-----------|------------|-----------|----------|--------------------|--------------|
| cluster-1 | 1          | true      | false    | unknown            | unknown      |
| cluster-2 | 2          | true      | false    | ok                 | ok           |

Islands:

| Island ID | Clusters             |
|-----------|----------------------|
| 1         | cluster-1, cluster-2 |

VPLEXcli:/> **cluster summary**

Clusters:

| Name      | Cluster ID | Connected | Expelled | Operational Status | Health State |
|-----------|------------|-----------|----------|--------------------|--------------|
| cluster-1 | 1          | false     | -        | -                  | -            |
| cluster-2 | 2          | true      | false    | degraded           | degraded     |

Islands:

| Island ID | Clusters  |
|-----------|-----------|
| 2         | cluster-2 |

Connectivity problems:

| From      | Problem   | To        |
|-----------|-----------|-----------|
| cluster-2 | can't see | cluster-1 |

VPLEXcli:/> **ll /clusters/cluster-1**

Attributes:

| Name                   | Value        |
|------------------------|--------------|
| allow-auto-join        | -            |
| auto-expel-count       | -            |
| auto-expel-period      | -            |
| auto-join-delay        | -            |
| cluster-id             | 7            |
| connected              | false        |
| default-cache-mode     | -            |
| default-caw-template   | true         |
| director-names         | [DirA, DirB] |
| island-id              | -            |
| operational-status     | not-running  |
| transition-indications | []           |
| transition-progress    | []           |
| health-state           | unknown      |
| health-indications     | []           |

- See also**
- ◆ [cluster add on page 77](#)
  - ◆ [cluster expel on page 83](#)
  - ◆ [cluster forget on page 85](#)

- ◆ [director shutdown on page 201](#)

## cluster status

Displays a cluster's operational status and health state.

**Contexts** All contexts.

**Syntax** cluster status

**Example** Show cluster status for a healthy cluster:

```
Vplexcli:/> cluster status
Cluster cluster-1
  operational-status:          ok
  transitioning-indications:
  transitioning-progress:
  health-state:                ok
  health-indications:
  local-com:                   ok

Cluster cluster-2
  operational-status:          ok
  transitioning-indications:
  transitioning-progress:
  health-state:                ok
  health-indications:
  local-com:                   ok

wan-com: ok
```

**Example** Show cluster status when a port group is not healthy:

```
Vplexcli:/> cluster status
Cluster cluster-1
  operational-status:          ok
  transitioning-indications:
  transitioning-progress:
  health-state:                ok
  health-indications:
  local-com:                   ok

Cluster cluster-2
  operational-status:          ok
  transitioning-indications:
  transitioning-progress:
  health-state:                ok
  health-indications:
local-com:                      connectivity: PARTIAL

port-group-1 - FAIL - No connectivity was found from any com port.
port-group-0 - OK - All expected connectivity is present.
```

**Example** Show cluster status during an inter-cluster link outage:

```
Vplexcli:/> cluster status
Cluster cluster-1
  operational-status:          degraded
  transitioning-indications:  suspended exports,suspended volumes
  transitioning-progress:
  health-state:                minor-failure
  health-indications:         169 suspended Devices
                              250 non-running remote virtual-volumes.
                              250 unhealthy Devices or storage-volumes
                              storage-volume unreachable
```

```

Cluster local-com: ok
cluster-2
operational-status: degraded
transitioning-indications: suspended exports,suspended volumes
transitioning-progress:
health-state: minor-failure
health-indications: 227 suspended Devices
                    250 non-running remote virtual-volumes.
                    250 unhealthy Devices or storage-volumes
                    storage-volume unreachable
local-com: ok

```

```
wan-com: connectivity: PARTIAL
```

```
port-group-1 - FAIL - No connectivity was found from any com port.
```

```
port-group-0 - OK - All expected connectivity is present.
```

**Example** Show cluster status when one cluster is shut down or expelled:

```
Vplexcli: /> cluster status
```

```

Cluster cluster-1
operational-status: not-running
transitioning-indications:
transitioning-progress:
health-state: unknown
health-indications:

Cluster cluster-2
operational-status: degraded
transitioning-indications: suspended exports,suspended volumes
transitioning-progress:
health-state: degraded
health-indications: 2 suspended Devices

```

Table 6 cluster status field descriptions

| Field                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| operational status         | <p>Operational status of the cluster. During transition periods cluster moves from one operational state to another.</p> <p><b>cluster departure</b> - One or more of the clusters cannot be contacted. Commands affecting distributed storage are refused.</p> <p><b>degraded</b> - The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volumes, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, out-of-date devices, and so forth.</p> <p><b>device initializing</b> - If clusters cannot communicate with each other, then the distributed-device will be unable to initialize.</p> <p><b>device out of date</b> - Child devices are being marked fully out of date. Sometimes this occurs after a link outage.</p> <p><b>expelled</b> - The cluster has been isolated from the island either manually (by an administrator) or automatically (by a system configuration setting).</p> <p><b>ok</b> - The cluster is operating normally.</p> <p><b>shutdown</b> - The cluster's directors are shutting down.</p> <p><b>suspended exports</b> - Some I/O is suspended. This could be result of a link failure or loss of a director. Other states might indicate the true problem.</p> <hr/> <p><b>Note:</b> it may <i>not</i> be a problem, and the VPLEX might be waiting for you to confirm the resumption of I/O.</p> <hr/> <p><b>transitioning</b> - Components of the software are recovering from a previous incident (for example, the loss of a director or the loss of an inter-cluster link).</p> |
| transitioning -indications | Additional information if the transitioning-progress is anything other than blank.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| transitioning-progress     | Indicates progress for supported transitions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| health-state               | <p><b>critical failure</b> - The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity.</p> <p><b>degraded</b> - The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volumes, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, or out-of-date devices.</p> <p><b>ok</b> - The cluster is functioning normally.</p> <p><b>unknown</b> - VPLEX cannot determine the cluster's health state, or the state is invalid.</p> <p><b>major failure</b> - The cluster is failing and some functionality may be degraded or unavailable. This may indicate complete loss of back-end connectivity.</p> <p><b>minor failure</b> - The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| health-indications         | Additional information if the health-state field is anything other than "ok".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| local-com                  | <p><b>ok</b> - All wan-com links have the expected connectivity: this port-group is operating correctly.</p> <p><b>warning</b> - Some links have unexpected connectivity. This port-group is operational but not properly configured. Performance may not be optimal.</p> <p><b>error</b> - Some connectivity is missing from this port-group. It is not operating correctly.</p> <p><b>fail</b> - All connectivity is missing from this port-group. wan-com is not operational.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| wan-com                    | <p><b>full</b> - All port-groups have a status of either ok or warning. wan-com connectivity is complete through minor configuration errors may still exist. See individual port-group statuses.</p> <p><b>partial</b> - Some port-groups have a status of error or fail, but at least one port-group has a status of ok or warning. wan-com is operating (possibly minimally) through at least one channel. Performance is degraded.</p> <p><b>none</b> - All port-groups have a status of either error or fail. wan-com is not operational.</p> <p><b>not-applicable</b> - The system is a single-cluster (i.e. Local) system. Validating wan-com connectivity is not applicable.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

- See also**
- ◆ [cluster summary on page 93](#)
  - ◆ [ds summary on page 234](#)

## cluster summary

Displays a summary of all clusters and the connectivity between them.

**Contexts** All contexts.

In /clusters context, command is **summary**.

**Syntax** cluster summary

**Example** Display summary for healthy clusters:

```
VPlexcli:/> cluster summary
Clusters:
  Name          Cluster ID  Connected  Expelled  Operational Status  Health State
  -----
  cluster-1     1           true       false     ok                   ok
  cluster-2     2           true       false     ok                   ok

Islands:
  Island ID  Clusters
  -----
  1          cluster-1, cluster-2
```

**Example** Display cluster summary for VPLEX Metro configuration with a inter-cluster link outage:

```
VPlexcli:/> cluster summary
Clusters:
  Name          Cluster ID  Connected  Expelled  Operational Status  Health State
  -----
  cluster-1     1           true       false     degraded            minor-failure
  cluster-2     2           true       false     degraded            minor-failure

Islands:
  Island ID  Clusters
  -----
  1          cluster-1
  2          cluster-2
```

```
Connectivity problems:
  From      Problem      To
  -----
  cluster-2  can't see    cluster-1
  cluster-1  can't see    cluster-2
```

**Example** Display cluster summary for VPLEX Metro configuration with a cluster expelled:

```
VPlexcli:/> cluster summary
Clusters:
  Name          Cluster ID  Connected  Expelled  Operational Status  Health State
  -----
  cluster-1     1           true       true      isolated            degraded
  cluster-2     2           true       false     degraded            degraded

Islands:
  Island ID  Clusters
  -----
  1          cluster-1
  2          cluster-2
```

**Table 7 cluster summary field descriptions**

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Clusters:</b>   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Name               | Name of the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Cluster ID         | Cluster ID. For VPLEX Local, always 1. For VPLEX Metro or Geo1 or 2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Connected          | Whether or not the CLI is connected to at least one director in the cluster (connected to the cluster).<br><b>true</b> - CLI is connected to the cluster.<br><b>false</b> - CLI is not connected to the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Expelled           | <b>true</b> - The cluster is expelled from its island.<br><b>false</b> -The cluster is not expelled from its island.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Operational Status | <b>degraded</b> - The cluster is not operating as configured and is not currently transitioning. Examples include: degraded redundancy level (a director is dead), all exports switched to write through because of HWH problems, suspended virtual volumes / exports, storage volumes not visible from all directors, meta-volume not yet processed.<br><b>isolated</b> - The cluster is not communicating with any other clusters.<br><b>ok</b> - The cluster is functioning normally.<br><b>transitioning</b> - The cluster is reacting to external events and may not be operating as configured. I/O may be suspended during the transition period.<br><br><b>Note:</b> If no meta-volume has been configured, operational status is "transitioning".<br><br><b>unknown</b> - The VPLEX encountered a problem determining the operational status of the cluster. This may indicate a degraded state, since it usually means that at least one of the directors is not responding or is communicating abnormally. |
| Health State       | <b>critical failure</b> - The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity.<br><b>degraded</b> - The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volumes, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, out-of-date devices, and so forth.<br><b>ok</b> - The cluster is functioning normally.<br><b>unknown</b> - The VPLEX cannot determine the cluster's health state, or the state is invalid.<br><b>major failure</b> - The cluster is failing and some functionality may be degraded or unavailable. This may indicate complete loss of back-end connectivity.<br><b>minor failure</b> - The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.                                                                                               |
| <b>Islands:</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Island ID          | ID of the island. For current release, always 1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Clusters           | Names of clusters belonging to the island. For current release, always cluster-1 or cluster-2.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

**See also** ♦ [cluster status on page 89](#)

## cluster unexpel

Allows a cluster to rejoin the VPLEX.

**Contexts** All contexts.  
In /clusters context, command is **unexpel**.

**Syntax** `cluster unexpel  
[-c|--cluster] context path`

**Arguments** **Required arguments**  
`[-c|--cluster] context path` - Cluster to unexpel.

**Description** Clears the expelled flag for the specified cluster, allowing it to rejoin the VPLEX.

**Example** To manually unexpel a cluster, do the following:

1. Use the **cluster summary** command to verify that the cluster is expelled.
2. Use the **ll** command in the target cluster's cluster context to display the cluster's **allow-auto-join** attribute setting.

If the cluster's **allow-auto-join** attribute is set to **true**, the cluster automatically rejoins the system. Skip to [step 5](#)

If the cluster's **allow-auto-join** flag was set to false, proceed to [step 3](#)

3. Navigate to the target cluster's cluster context and use the **set** command to set the cluster's **allow-auto-join** flag to true. For example:

```
Vplexcli:/ cd clusters/cluster-2
Vplexcli:/clusters/cluster-2> set cluster-1::allow-auto-join true
```

4. Use the **cluster unexpel** command in target cluster's cluster context to manually unexpel a cluster, allowing the cluster to rejoin VPLEX. The syntax for the command is:

```
cluster unexpel cluster-id>
```

For example:

```
Vplexcli:/clusters> cluster unexpel --cluster cluster-2
```

5. Use the **cluster summary** command to verify all clusters are in one island and working as expected.

```
Vplexcli:/clusters> cluster summary
```

Clusters:

| Name      | Cluster ID | Connected | Expelled | Operational Status | Health State |
|-----------|------------|-----------|----------|--------------------|--------------|
| cluster-1 | 1          | true      | false    | ok                 | ok           |
| cluster-2 | 2          | true      | false    | ok                 | ok           |

Islands:

| Island ID | Clusters             |
|-----------|----------------------|
| 1         | cluster-1, cluster-2 |

**See also** ♦ [cluster expel on page 83](#)

## cluster-witness configure

Creates the cluster-witness context for enabling VPLEX Witness functionality and configuration commands.

**Contexts** All contexts.

**Syntax** `cluster-witness configure  
[-i|--ip-address] public-ip-address`

**Arguments** **Required arguments**

`[-i|--ip-address] public-ip-address` - \* Public IP address of the Cluster Witness server used as an end point of the IPsec tunnel and for diagnostics.

**Note:** This public ip-address should be the same ip-address provided to the “[configuration cw-vpn-configure](#)” command.

\* - argument is positional.

**Description** Cluster Witness is an optional component of VPLEX Metro and VPLEX Geo configurations.

Cluster Witness monitors both clusters and updates the clusters with its guidance, when necessary. Cluster Witness allows VPLEX to distinguish between inter-cluster link failures versus cluster failures, and to apply the appropriate detach-rules and recovery policies.



### IMPORTANT

**This command must be run on both management servers to create cluster-witness CLI contexts on the VPLEX.**



### IMPORTANT

**ICMP traffic must be permitted between clusters for this command to work properly.**

To verify that ICMP is enabled, log in to the shell on the management server and use the `ping IP address` command where the IP address is for a director in the VPLEX.

If ICMP is enabled on the specified director, a series of lines is displayed:

```
service@ManagementServer:~> ping 128.221.252.36
PING 128.221.252.36 (128.221.252.36) 56(84) bytes of data.
64 bytes from 128.221.252.36: icmp_seq=1 ttl=63 time=0.638 ms
64 bytes from 128.221.252.36: icmp_seq=2 ttl=63 time=0.591 ms
64 bytes from 128.221.252.36: icmp_seq=3 ttl=63 time=0.495 ms
64 bytes from 128.221.252.36: icmp_seq=4 ttl=63 time=0.401 ms
64 bytes from 128.221.252.36: icmp_seq=5 ttl=63 time=0.552 ms

--- 128.221.252.36 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4002ms
rtt min/avg/max/mdev = 0.401/0.535/0.638/0.084 ms
```

If ICMP is disabled, nothing is displayed.

Press Ctrl-C to exit from ping.

**Example** In the following example:

- ◆ The **ls** command at the root context shows that the Cluster Witness context is not configured.
- ◆ The **cluster-witness configure** command creates the cluster-witness context.
- ◆ The **ls** command shows the Public and Private IP addresses and the administrative Cluster Witness state.

```
Vplexcli:/> ls
```

```
clusters/                data-migrations/
distributed-storage/     engines/                management-server/
monitoring/              notifications/          system-defaults/
```

```
Vplexcli:/> cluster-witness configure -i 10.31.25.45
```

```
Vplexcli:> ls /cluster-witness
```

```
Attributes:
```

| Name               | Value         |
|--------------------|---------------|
| -----              | -----         |
| admin-state        | disabled      |
| private-ip-address | 128.221.254.3 |
| public-ip-address  | 10.31.25.45   |

```
Contexts:
```

```
components
```

- See also**
- ◆ *VPLEX Procedure Generator - VPLEX Witness: Install and Setup*
  - ◆ [cluster-witness disable on page 98](#)
  - ◆ [cluster-witness enable on page 101](#)
  - ◆ [configuration cw-vpn-configure on page 115](#)

## cluster-witness disable

Disables Cluster Witness on both management servers and on Cluster Witness Server.

**Contexts** All contexts.

In /cluster-witness context, command is **disable**.

**Note:** This command is available only after Cluster Witness has been configured and cluster-witness CLI context is visible.

**Syntax**

```
cluster-witness disable
  [-f|--force]
  [-w|--force-without-server]
```

**Arguments** **Optional arguments**

**[-f|--force]** - Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.

**[-w|--force-without-server]** - Force the operation to disable Cluster Witness on both clusters when connectivity to Cluster Witness Server is lost but the two clusters are connected.

Use this option when Cluster Witness fails or disconnects from both clusters and recovery is unlikely to happen soon.



### CAUTION

**Use the --force-without-server option with extreme care. Use this option to disable Cluster Witness in order to use configured rule-sets for I/O to distributed volumes in consistency groups.**

**Note:** If Cluster Witness Server becomes reachable when --force-without-server option is used, the command will also disable the Cluster Witness Server.

**Description** Disables Cluster Witness on both management servers and on Cluster Witness Server. Allows consistency group rule-sets to dictate I/O behavior to distributed virtual volumes in consistency groups.



### IMPORTANT

**Cluster Witness has no effect on distributed virtual volumes outside of consistency groups.**



### CAUTION

**Use this command from only one management server.**

Disabling Cluster Witness does not imply that Cluster Witness components are shut down. If Cluster Witness is disabled, the clusters stop sending health-check traffic to the Cluster Witness Server and the Cluster Witness Server stops providing guidance back to the clusters.

**Note:** If the Cluster Witness Server or connectivity to the Cluster Witness Server will be not operational for a long period, use the **--force-without-server** argument. This prevents a system-wide Data Unavailability of all distributed virtual volumes in consistency groups if an

additional inter-cluster link communication or cluster failure occurs while there is no access to Cluster Witness Server *and* Cluster Witness is enabled. Once Cluster Witness Server is accessible from both management servers, use the “`cluster-witness enable`” command to re-enable the functionality.

Automatic pre-checks ensure that the Cluster Witness configuration is in a state where it can be disabled. Pre-checks:

- ◆ Verify management connectivity between the management servers
- ◆ Verify connectivity between management servers and the Cluster Witness Server
- ◆ Verify all the directors are up and running

**Note:** If the `--force-without-server` option is used, the automatic pre-check to verify connectivity between management servers and the Cluster Witness Server is not performed.

- ◆ Verify connectivity between directors and each management server
- ◆ Verify that Cluster Witness is configured on both clusters
- ◆ Verify that the metavolume on both clusters is healthy

**Example** Disable Cluster Witness from the root context:

```
VPlexcli:/> cluster-witness disable
```

```
"WARNING: Disabling Cluster Witness may cause data unavailability in the event of a disaster.
Please consult the VPLEX documentation to confirm that you would like to disable Cluster
Witness. Continue? Yes
```

**Example** Disable Cluster Witness from the cluster-witness context when the Cluster Witness Server is not reachable. In the following example:

- ◆ The `disable` command fails because the Cluster Witness Server is not reachable.
- ◆ The `disable --force-without-server` command disables Cluster Witness.
- ◆ The `ll /components` command displays the state of the Cluster Witness configuration.

```
VPlexcli:/cluster-witness> disable
```

```
WARNING: Disabling Cluster Witness may cause data unavailability in the event of a disaster.
Please consult the VPLEX documentation to confirm that you would like to disable Cluster
Witness. Continue? (Yes/No) y
```

```
cluster-witness disable: Evaluation of <<disable>> failed.
cause: Could not disable Cluster Witness.
cause: Cluster Witness cannot be disabled due to failure of a pre-check.
cause: Unable to communicate with Cluster Witness Server. Please check the
state of the Cluster Witness Server and its
connectivity and try again
```

```
VPlexcli:/cluster-witness> disable --force-without-server
```

```
WARNING: Disabling Cluster Witness may cause data unavailability in the event of a disaster.
Please consult the VPLEX documentation to confirm that you would like to disable Cluster
Witness. Continue? (Yes/No) y
```

```
VPlexcli:/cluster-witness> ll components/
```

```
/cluster-witness/components:
Name          ID  Admin State  Operational State  Mgmt Connectivity
-----
cluster-1     1  disabled    -                   ok
cluster-2     2  disabled    -                   ok
server        -  unknown     -                   failed
```

- See also**
- ◆ *VPLEX Procedure Generator - VPLEX Witness: Install and Setup*
  - ◆ [cluster summary on page 93](#)
  - ◆ [cluster-witness configure on page 96](#)
  - ◆ [cluster-witness enable on page 101](#)
  - ◆ [configuration cw-vpn-configure on page 115](#)
  - ◆ [connectivity validate-wan-com on page 145](#)
  - ◆ [vpn status on page 514](#)

## cluster-witness enable

Enables Cluster Witness on both clusters and Cluster Witness Server in a VPLEX Metro or Geo configuration.

**Contexts** All contexts.

In /cluster-witness context, command is **enable**.

**Note:** This command is available only after Cluster Witness has been configured and cluster-witness CLI context is visible.

**Syntax** `cluster-witness enable`

### Description



#### CAUTION

**Use this command from the management server on only one cluster.**

Automatic pre-checks run before the **cluster-witness enable** command is issued. Pre-checks verify that the VPLEX is in a state that Cluster Witness can be enabled. Pre-checks:

- ◆ Verify management connectivity between both the management servers
- ◆ Verify connectivity between each management server and the Cluster Witness Server
- ◆ Verify connectivity between directors and each management server
- ◆ Verify that Cluster Witness CLI context is configured on both clusters
- ◆ Verify that a metavolume is present and healthy on both clusters
- ◆ Verify all the directors are healthy

If any of the pre-checks fail, the command displays the cause of failure on the specific component and warns about possible Data Unavailability risk, if any.



#### WARNING

*There is no rollback mechanism. If the enable command fails on some components and succeeds on others, it may leave the system in an inconsistent state. If this occurs, consult the Troubleshooting Guide and/or contact EMC Customer Support.*

The cluster-witness context does not appear in the VPLEX CLI unless the context has been created using the cluster-witness configure command. The cluster-witness CLI context appears under the root context. The cluster-witness context includes the following sub-contexts:

- ◆ /cluster-witness/components/cluster-1
- ◆ /cluster-witness/components/cluster-2
- ◆ /cluster-witness/components/server

**Example** In the following example:

- ◆ The **ll** command verifies that Cluster Witness is configured (the context exists)
- ◆ The **cd** command changes the context to cluster-witness
- ◆ The **cluster-witness enable** command enables Cluster Witness

- ◆ The **ll /components/\*** command displays the components on cluster-1, cluster-2, and the Cluster Witness Server:

```
Vplexcli:/> ll /cluster-witness
```

```
Attributes:
```

| Name               | Value         |
|--------------------|---------------|
| admin-state        | disabled      |
| private-ip-address | 128.221.254.3 |
| public-ip-address  | 10.31.25.235  |

```
Contexts:
```

| Name       | Description                |
|------------|----------------------------|
| components | Cluster Witness Components |

```
Vplexcli:/> cd /cluster-witness
```

```
Vplexcli:/cluster-witness> cluster-witness enable
```

```
Vplexcli:/cluster-witness> ll /components/*
```

```
/cluster-witness/components/cluster-1:
```

| Name                    | Value                                                                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| admin-state             | enabled                                                                                                                                |
| diagnostic              | INFO: Current state of cluster-1 is in-contact (last state change: 0 days, 56 secs ago; last message from server: 0 days, 0 secs ago.) |
| id                      | 1                                                                                                                                      |
| management-connectivity | ok                                                                                                                                     |
| operational-state       | in-contact                                                                                                                             |

```
/cluster-witness/components/cluster-2:
```

| Name                    | Value                                                                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| admin-state             | enabled                                                                                                                                |
| diagnostic              | INFO: Current state of cluster-2 is in-contact (last state change: 0 days, 56 secs ago; last message from server: 0 days, 0 secs ago.) |
| id                      | 2                                                                                                                                      |
| management-connectivity | ok                                                                                                                                     |
| operational-state       | in-contact                                                                                                                             |

```
/cluster-witness/components/server:
```

| Name                    | Value                                                                                                                                                                                                                  |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| admin-state             | enabled                                                                                                                                                                                                                |
| diagnostic              | INFO: Current state is clusters-in-contact (last state change: 0 days, 56 secs ago.) (last time of communication with cluster-2: 0 days, 0 secs ago.) (last time of communication with cluster-1: 0 days, 0 secs ago.) |
| id                      | -                                                                                                                                                                                                                      |
| management-connectivity | ok                                                                                                                                                                                                                     |
| operational-state       | clusters-in-contact                                                                                                                                                                                                    |

Table 8 cluster witness display fields (1 of 2)

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Name               | Name of component,<br>For VPLEX clusters – name assigned to cluster.<br>For VPLEX Witness server – “server”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| id                 | ID of a VPLEX cluster.<br>Always blank “-” for Witness server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| admin state        | Identifies whether VPLEX Witness is enabled/disabled. Valid values are:<br><b>enabled</b> - VPLEX Witness functionality is enabled on this component.<br><b>disabled</b> - VPLEX Witness functionality is disabled on this component.<br><b>inconsistent</b> - All Cluster Witness components are reachable over the management network but some components report their administrative state as disabled while others report it as enabled. This is a rare state which may result failure during enabling or disabling.<br><b>unknown</b> - This component is not reachable and its administrative state cannot be determined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| private-ip-address | Private IP address of the Cluster Witness Server VM used for cluster witness-specific traffic.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| public-ip-address  | Public IP address of the Cluster Witness Server VM used as an endpoint of the IPsec tunnel.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| diagnostic         | String generated by CLI based on the analysis of the data and state information reported by the corresponding component.<br><b>WARNING: Cannot establish connectivity with Cluster Witness Server to query diagnostic information.</b> - Cluster Witness Server or one of the clusters is unreachable.<br><b>Local cluster-x hasn't yet established connectivity with the server</b> - The cluster has never connected to Cluster Witness Server.<br><b>Remote cluster-x hasn't yet established connectivity with the server</b> - The cluster has never connected to Cluster Witness Server.<br><b>Cluster-x has been out of touch from the server for X days, Y secs</b> - Cluster Witness Server has not received messages from a given cluster for longer than 60 seconds.<br><b>Cluster witness server has been out of touch for X days, Y secs</b> - Either cluster has not received messages from Cluster Witness Server for longer than 60 seconds.<br><b>Cluster Witness is not enabled on component-X, so no diagnostic information is available</b> - Cluster Witness Server or either of the clusters is disabled. |

**Table 8 cluster witness display fields (2 of 2)**

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operational State | <p>Operational state of the corresponding component.</p> <p>For clusters – Healthy state is in-contact. All other states indicate a problem.</p> <p>For Cluster Witness Server – Healthy state is clusters-in-contact. All other states indicate a problem.</p> <p><u>For Cluster Witness Server:</u></p> <p><b>clusters-in-contact</b> - Both clusters are in contact with each other over the inter-cluster network.</p> <p><b>cluster-partition</b> - The clusters are partitioned from each other over the inter-cluster network, and the Cluster Witness Server can still talk to each cluster. See NOTE: below.</p> <p><b>cluster-partition</b> - The clusters are partitioned from each other over the inter-cluster network, and the Cluster Witness Server can still talk to each cluster. See NOTE: below.</p> <p><b>cluster-unreachable</b> - One cluster has either failed or become isolated (that is partitioned from its peer cluster and disconnected from the Cluster Witness Server). See NOTE: below.</p> <p><b>unknown</b> - Cluster Witness Server does not know the states of one or both of the clusters and needs to learn them before it can start making decisions. Cluster Witness Server assumes this state upon startup.</p> <hr/> <p><b>Note:</b> When the server's operational state is cluster-partition or cluster-unreachable, this may not reflect the current observation of the Cluster Witness Server. The Cluster Witness Server state and the guidance that it provides to the clusters based on its state is persistent -- if Cluster Witness Server observes a failure (changes its state and provides guidance to the clusters), the Cluster Witness Server maintains this state even if current observations change. Cluster Witness Server maintains its failure state and guidance until both clusters and their connectivity fully recover. This policy is implemented in order to avoid potential data corruption scenarios due to split brain (that would be possible if this semantics were not followed).</p> <hr/> <p><u>For VPLEX clusters:</u></p> <p><b>in-contact</b> - This cluster is in contact with its peer over the inter-cluster network. Rebuilds may be in progress. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups is allowed from the perspective of VPLEX Witness.</p> <p><b>cluster-partition</b> - This cluster is not in contact with its peer and the Cluster Witness Server has declared that two clusters partitioned. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of VPLEX Witness.</p> <p><b>remote-cluster-isolated-or-dead</b> - This cluster is not in contact with its peer and the Cluster Witness Server has declared that the remote cluster (the peer) was isolated or dead. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of VPLEX Witness.</p> <p><b>local-cluster-isolated</b> - This cluster is not in contact with its peer and the Cluster Witness Server has declared that the remote cluster (the peer) is the only proceeding cluster. This cluster must suspend I/O to all distributed virtual volumes in all synchronous consistency groups regardless of cluster preference.</p> <p>NOTE: When a cluster is isolated from both the remote cluster and Cluster Witness Server, its state is unknown. When connectivity to VPLEX Witness server is restored, the state of this cluster changes to local-cluster-isolated if this cluster remains partitioned from the peer cluster.</p> <p><b>unknown</b> - This cluster is not in contact with its peer over the inter-cluster network and is awaiting guidance from the Cluster Witness Server. I/O to all distributed virtual volumes in all synchronous consistency groups is suspended regardless of cluster preference.</p> <p>NOTE: When its state is local-cluster-isolated, the cluster does not receive guidance from Cluster Witness Server. If connectivity with Cluster Witness Server is restored before the inter-cluster link is restored state is local-cluster-isolated.</p> |
| Mgmt Connectivity | <p>Reachability of the specified Witness component over the IP management network from the management server where the CLI command is run.</p> <p><b>ok</b> - The component is reachable</p> <p><b>failed</b> - The component is not reachable</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

- See also**
- ◆ [VPLEX Procedure Generator - VPLEX Witness: Install and Setup](#)
  - ◆ [cluster summary on page 93](#)
  - ◆ [cluster-witness configure on page 96](#)
  - ◆ [cluster-witness disable on page 98](#)

- ◆ [configuration cw-vpn-configure](#) on page 115
- ◆ [vpn status](#) on page 514

## collect-diagnostics

Collects the two latest core files from each component, logs, and configuration information from the management server and directors.

**Contexts** All contexts.

**Syntax**

```
collect-diagnostics
  --notrace
  --nocores
  --noperf
  --noheap
  --noextended
  --faster
  --local-only
  --minimum
  --allcores
  --large-config
  --recoverpoint-only
  --out-dir directory
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**--notrace** - Do not collect fast trace dump files from the directors.

**--nocores** - Do not collect core files from the directors.

**--noperf** - Do not collect performance sink files.

**--noheap** - Do not dump the management console's heap.

**--noextended** - Omit collection of extended diagnostics. Implies use of **--nocores**, **--noheap**, and **--notrace** arguments.

**--faster** - Omits some of the more time-consuming operations. Use only when **collect-diagnostics** is expected to take very long, for example on large configurations. The following commands are omitted: **getsysinfo**, **export port summary** and **connectivity director**.

**--local-only** - Gathers diagnostics only from the local cluster and directors.

**--minimum** - Combines all space-saving and time-saving options and operates only on the local cluster.

Use this argument when the **collect-diagnostics** command is expected to take a long time or produce excessively large output files. Combines the **--noextended**, **--faster** and **--local-only** arguments. Also omits the second **sms dump** command and output file (**smsDump-CLI\_Logs\_<timestamp>.zip**).

**--allcores** - Collect all available core files from the directors. By default, only the two latest core files are collected.

**--large-config** - Omits the cluster **configdump** command output. Use this argument only when the configuration has:

- ◆ 8000 or more storage volumes, and
- ◆ 4000 or more local top-level devices, and
- ◆ 2000 or more distributed devices.

**--recoverpoint-only** - Collects only RecoverPoint diagnostic information.

**--out-dir *directory*** - The directory into which to save the zip file containing the collected diagnostics.

Default: /diag/collect-diagnostics-out

Not needed for normal usage.

## Description

Collects logs, cores and configuration information from management server and directors. Places the collected output files in the /diag/collect-diagnostics-out directory on the management server.

2 compressed files are placed in /diag/collect-diagnostics-out directory:

- ◆ **fla-diagnostics-extended-timestamp.tar.gz** - Contains java heap dump, fast trace dump, two latest core files (if they exist).
- ◆ **fla-diagnostics-timestamp.tar.gz** - Contains everything else.

Best practice is to collect both files. The extended collect-diagnostics file is usually large, and thus takes some time to transfer.

Recommended practice is to transfer the base collect-diagnostics file (**fla-diagnostics-timestamp.tar.gz**) first and begin analysis while waiting for the extended file to transfer.

Starting in GeoSynchrony 5.1, this command also collects RecoverPoint diagnostics data:

- ◆ The extended .tar file includes the two latest RecoverPoint kdriver core files.
- ◆ The base tar file includes an additional directory: /opt/recoverpoint/. Files in this directory include the RecoverPoint splitter logs:
- ◆ A zip file named **director-name-time-stamp.zip** contains the following splitter logs:
  - vpsplitter.log.xx -
  - vpsplitter.log.periodic\_env -
  - vpsplitter.log.current\_env. -

**Note:** On VPLEX Metro or Geo configurations, collect-diagnostics must be invoked from each management server in order to collect complete diagnostics. Director core files and management server diagnostics files are only collected for the local cluster.

Other than core files, director diagnostics are retrieved from ALL directors in a VPLEX Metro or Geo unless the **--local-only** argument is used.

**Note:** Core files are always collected only from local directors. Starting in Release 5.1, only the latest 2 core files are collected by default, and any older core files are not collected. To collect all the core files, use the **--allcores** argument.



## CAUTION

In VPLEX Metro and Geo configurations, run the **collect-diagnostics** command on each management server, but NOT at the same time. Even if the **--local-only** argument is used, do not run the command on both management servers at the same time.

EMC recommends that files created by **collect-diagnostics** be removed from the management server as soon as possible to avoid filling management server disk partitions.

**Example** Collect diagnostics, omitting the core files on the directors and the management server console heap, and send the output to the default directory:

```
VPllexcli:/> collect-diagnostics --nocores --noheap
```

**Example** Collect all RecoverPoint diagnostics, including all available RecoverPoint core files, and send the output to the default directory:

```
VPllexcli:/> collect-diagnostics --recoverpoint-only --allcores
```

**See also**

- ◆ [cluster configdump on page 81](#)
- ◆ [director appdump on page 190](#)
- ◆ [getsysinfo on page 272](#)
- ◆ [sms dump on page 448](#)

## configuration complete-system-setup

Completes the VPLEX Metro or Geo configuration.

**Contexts** All contexts.

**Syntax** `configuration complete-system-setup`

**Description** Completes the automated EZ-Setup Wizard for VPLEX Metro and VPLEX Geo configurations.

This command must be run twice: once on each cluster.

**Note:** Before using this command on either cluster, first use the “[configuration system-setup](#)” command (on both clusters).

**Example** `Vplexcli:/> configuration complete-system-setup`

**See also**

- ◆ [configuration connect-remote-directors](#) on page 113
- ◆ [configuration continue-system-setup](#) on page 114
- ◆ [configuration system-reset](#) on page 134
- ◆ [configuration system-setup](#) on page 135

## configuration configure-auth-service

Configures the authentication service selected by the user.

**Contexts** All contexts.

**Syntax** configuration configure-auth-service

**Description** Configures the selected authentication service.

See [authentication directory-service configure on page 34](#) for a description of the available authentication services.

**Example** Configure the selected authentication service:

```
VPllexcli:/> configuration configure-auth-service
```

```
Configure Authentication Service Provider (Optional)
```

You may select to use your existing LDAP or Active Directory as a directory service to authenticate VPLEX users. To configure this, you will need the authentication service provider server information, and the security information to map the users.

Or, you may choose not to configure an authentication service provider at this time. You may configure an authentication service provider for authentication at any time, using VPLEX CLI commands.

```
Would you like to configure an authentication service provider to authenticate VPLEX users?
(yes/no) [no]:yes
```

VPLEX supports the following types of authentication service providers:

1. LDAP
2. AD

```
Select the type of authentication service provider you would like use for VPLEX authentication.
(1 - 2) [1]: 1
```

```
Enter the Authentication Server IP Address: []: 10.31.52.53
```

VPLEX supports these connections types:

1. SSL
2. TLS

```
Select your connection type (1 - 2) [2]:
```

```
Enter the port to be used for LDAP [389]:
```

Configure Security Settings for Your Authentication Service Provider

To configure the Authentication Service Provider you will need: the base distinguished name, the bind distinguished name, and the mapprincipal. Examples of these are:

Base Distinguished Name Example: dc=security,dc=orgName,dc=companyName,dc=com

Bind Distinguished Name Example: cn=Administrator,dc=security,dc=orgName,dc=companyName,dc=com

Mapprincipal Example: ou=people,dc=security,dc=orgName,dc=companyName,dc=com

```
Enter the Base Distinguished Name []: dc=security,dc=sve,dc=emc,dc=com
```

```
Enter the Distinguished Bind Name []: cn=Administrator,dc=security,dc=sve,dc=emc,dc=com
```

Enter the mappprincipal []: **ou=people,dc=security,dc=sve,dc=emc,dc=com**

cstadmin: Object PAM VPLEX-PAM-Authority configured.

Enter Administrator's password:

Connecting to authentication server (may take 3 minutes) ...  
Configuration of the ldap authentication service is complete

---

**Note:** After running this command, run the **webserver restart** command.

---

- See also**
- ◆ [authentication directory-service configure on page 34](#)
  - ◆ [authentication directory-service unconfigure on page 40](#)
  - ◆ [configuration system-setup on page 135](#)

## configuration connect-local-directors

Connects to the directors in the local cluster.

**Contexts** All contexts.

**Syntax** `configuration connect-local-directors`  
`[-f|--force]`

**Arguments** **Required arguments**  
 None.

**Optional arguments**  
`[-f|--force]` - Connect to local directors regardless of current connections.

**Description** This command executes **connect** commands to all local directors.  
 Use the **--force** argument if one or more local directors are already connected.  
 The connections use the director's default name. For example: director-1-1-A.

**Example** Connect the local directors to the cluster:

```
Vplexcli:/> configuration connect-local-directors
```

**Example** Use the **--force** argument when the directors are already connected:

```
Vplexcli:/> configuration connect-local-directors --force  

Already connected to Plex firmware director-1-1-A  

<128.221.252.35,128.221.253.35>.  
  

Already connected to Plex firmware director-1-1-B  

<128.221.252.36,128.221.253.36>.
```

**See also**

- ◆ [configuration continue-system-setup on page 114](#)
- ◆ [configuration connect-remote-directors on page 113](#)
- ◆ [configuration system-setup on page 135](#)
- ◆ [connect on page 137](#)

## configuration connect-remote-directors

Connects the local cluster to directors in the remote cluster.

**Contexts** All contexts.

**Syntax** `configuration connect-remote-directors`  
`[-f|--force]`

**Arguments** **Required arguments**  
None.

**Optional arguments**  
`[-f|--force]` - Connect to remote directors regardless of current connections.

**Description** During system setup for a VPLEX Metro or Geo configuration, use the **configuration connect-remote-directors** command to connect the local cluster to the directors in the remote cluster.

Run this command twice: once from the local cluster to connect to remote directors, and once from the remote cluster to connect to local directors.

Prerequisite: Number of directors at each cluster.

**Example** `Vplexcli:/> configuration connect-remote-directors`

**See also**

- ◆ [configuration continue-system-setup on page 114](#)
- ◆ [configuration connect-local-directors on page 112](#)
- ◆ [configuration system-setup on page 135](#)
- ◆ [connect on page 137](#)

## configuration continue-system-setup

Continues the EZ-Setup Wizard after back-end storage is configured and allocated for the cluster.

**Contexts** All contexts.

**Syntax** `configuration continue-system-setup`

**Description** This command validates the back-end configuration for the local cluster. The cluster must have its back-end allocated and configured for this command to succeed.

Use the “[configuration system-setup](#)” command to start the EZ-Setup Wizard to configure the VPLEX.

Zone the back-end storage to the port WWNs of the VPLEX back-end ports.

After the back-end storage is configured and allocated for the cluster, use the this command to complete the initial configuration.

**Example** `VPllexcli:/> configuration continue-system-setup`

**See also** ♦ [configuration system-setup on page 135](#)

## configuration cw-vpn-configure

Establishes VPN connectivity between a VPLEX management server and the Cluster Witness Server and starts the VPN tunnel between them. The command is interactive and requires inputs to complete successfully.

**This command must be run on both management servers to establish a 3-way VPN between the management servers and the Cluster Witness Server.**

**Contexts** All contexts.

**Syntax** `configuration cw-vpn-configure  
-i public-ip-address`

**Arguments** **Required arguments**  
[**-i** | **--ip-address**] *public-ip-address* - \* Valid dot-separated public IP address of the Cluster Witness Server.

### Description



#### IMPORTANT

**In order for this command to succeed, the following conditions must be true:**

- ◆ VPN connectivity is established between the VPLEX management servers
- ◆ Cluster Witness Server is successfully deployed based on the steps in the Cluster Witness Installation guide
- ◆ Cluster Witness Server is configured with a static public IP interface

This command:

- ◆ Configures the VPN tunnel between the local VPLEX management server and the Cluster Witness Server.
- ◆ Checks if VPLEX cluster's management servers are connected by the VPN tunnel. If Cluster Witness Server is not yet configured with VPN:
  - Generates the Cluster Witness Server host certificate
  - Configures the IPsec settings for Cluster Witness Server on the VPLEX management server
  - Configures the IPsec settings for the management server on the Cluster Witness Server
  - Restarts the IPsec service on the Cluster Witness Server VM
- ◆ Validates Cluster Witness Server VPN connectivity with the VPLEX management server.

#### **Prerequisites**

The following information is required to complete this command:

- ◆ VPLEX Metro or Geo setup is successfully completed using EZSetup. This creates a VPN connection between VPLEX management servers.
- ◆ Confirm that the VPLEX setup meets the requirement for Cluster Witness configuration. See the *VPLEX Cluster Witness Deployment and Configuration* guide for more information.
- ◆ Passphrase of the Certificate Authority Key that was provided during VPN configuration between management servers.

- ◆ Passphrase for creating the Cluster Witness host certificate. This passphrase should be different than passphrase used in creating host certificates for management servers.

**Note:** If this command is run when the VPN is already configured, the following error message is displayed: VPN connectivity is already established.

**Example** In the following example:

- ◆ Configure the VPN tunnel between VPLEX management server (10.31.25.26) to the Cluster Witness Server (10.31.25.45) while maintaining the pre-configured VPN connection to the other management server (10.31.25.27).
- ◆ Configure the VPN connection between VPLEX management server (10.31.25.27) to the Cluster Witness Server (10.31.25.45) while maintaining the pre-configured VPN connection to the other management server (10.31.25.26).

**Note:** The passphrase entered for the Cluster Witness Server Host Certificate Key is reconfirmed with a Re-enter the passphrase for the Certificate Key prompt. This is because the Cluster Witness Server Host Certificate Key password is created for the first time. There is no reconfirmation prompt for the passphrase of Certificate Authority because this CA passphrase Key is already created earlier as a part of complete-metro-setup.

In this example, run the command on the management server in the first VPLEX cluster (10.31.25.26):

```
VPlxcli:/> configuration cw-vpn-configure -i 10.31.25.45
```

```
The Cluster Witness requires a VPLEX Metro or VPLEX Geo configuration. Is this system
configured as a Metro or Geo? (Y/N): y
```

**Note:** If this is not an initial install, the following prompt appears:

“Please enter the passphrase for the Certificate Authority that was provided while configuring VPN between management servers.”

Enter the CA passphrase.

If this is an initial installation of Release 5.0 or later, the CA passphrase is provided automatically, and this question does not appear.

```
Enter the passphrase to create the Cluster Witness Server Host Certificate Key (at least 8
characters) :passphrase
```

```
Re-enter the passphrase for the Certificate Key: passphrase
```

```
New Host certificate request /etc/ipsec.d/reqs/cwsHostCertReq.pem created
```

```
New Host certificate /etc/ipsec.d/certs/cwsHostCert.pem created and signed by the CA
Certificate /etc/ipsec.d/cacerts/strongswanCert.pem
```

```
VPN Host Certificate created successfully
```

```
Please enter the IP address of the remote cluster management server that will be included in
the 3-way VPN setup: 10.31.25.27
```

```
Verifying the VPN status between the management servers...
```

```
IPSEC is UP
```

```
Remote Management Server at IP Address 10.31.25.27 is reachable
```

```
Remote Internal Gateway addresses are reachable
```

```
Verifying the VPN status between the management server and the cluster witness server...
```

```
Cluster Witness Server at IP Address 128.221.254.3 is not reachable
```

```
Verifying the VPN status between the management server and the cluster witness server...
```

```
IPSEC is UP
Cluster Witness Server at IP Address 128.221.254.3 is reachable
```

```
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.31.25.27 is reachable
Remote Internal Gateway addresses are reachable
```

The VPN configuration between this cluster and the Witness server is complete.

The Setup Wizard has completed the automated portion of configuring your cluster. From this point, please follow the manual procedures defined in the Installation and Setup Guide.

```
Vplexcli:/>
```

Run the command on the management server in the second VPLEX cluster (10.31.25.27):

```
Vplexcli:/> configuration cw-vpn-configure -i 10.31.25.45
```

```
Cluster witness requires a Vplex Metro or Vplex Geo configuration. Is this system a Metro or Geo? (Y/N): y
```

```
Please enter the IP address of the remote cluster management server in the plex that will be involved in the 3 way VPN setup: 10.31.25.26
```

```
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.31.25.26 is reachable
Remote Internal Gateway addresses are reachable
```

```
Verifying the VPN status between the management server and the cluster witness server...
Cluster Witness Server at IP Address 128.221.254.3 is not reachable
```

```
Verifying the VPN status between the management server and the cluster witness server...
IPSEC is UP
Cluster Witness Server at IP Address 128.221.254.3 is reachable
```

```
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.31.25.26 is reachable
Remote Internal Gateway addresses are reachable
```

Cluster Witness Server VPN Configuration Complete.

The Setup Wizard has completed the automated portion of configuring your cluster. From this point, please follow the manual procedures defined in the Installation and Setup Guide.

The log summary for configuration automation has been captured in /var/log/Vplex/cli/Vplexconfig.log

The task summary and the commands executed for each automation task has been captured in /var/log/Vplex/cli/Vplexcommands.txt

The output for configuration automation has been captured in /var/log/Vplex/cli/capture/Vplexconfiguration-session.txt

- See also**
- ◆ [VPLEX Procedure Generator.](#)
  - ◆ [cluster-witness configure on page 96](#)
  - ◆ [cluster-witness disable on page 98](#)
  - ◆ [cluster-witness enable on page 101](#)
  - ◆ [configuration system-setup on page 135](#)

## configuration cw-vpn-reset

Resets the VPN connectivity between the management server and the Cluster Witness Server.

**Contexts** All contexts.

**Syntax** configuration cw-vpn-reset

**Description** Resets the VPN between the management server on a cluster and the Cluster Witness Server.



### **WARNING**

*Use this command with EXTREME CARE. This command will erase all Cluster Witness VPN configuration.*

In order to complete, this command requires VPN connectivity between the management server and the Cluster Witness Server.

**Note:** Run this command twice: once from each management server.

**Example** From the first cluster:

```
Vplexcli:/> configuration cw-vpn-reset
```

```
This command will reset the vpn connectivity from this VPLEX instance to the Cluster Witness Server.
```

```
Do you want to continue? (Y/N): Y
```

```
To disable the VPN connectivity to the Cluster Witness Server please enter RESET (case sensitive): RESET Verifying if there
```

```
is a VPN connection between the Management Server and the Cluster Witness Server...
```

```
Verifying if the Cluster Witness has been configured on this Management Server...
```

```
Verifying if the Cluster Witness has been enabled on this Management Server...
```

```
Successfully removed the connection name and updated the Cluster Witness Server ipsec.conf file
Successfully transferred the ipsec configuration file to the Cluster Witness Server and
restarted the IPSec process
```

```
Successfully removed the cluster witness connection name from the Management Server ipsec.conf
file
```

```
Successfully restarted the ipsec process on the Management Server
```

```
Resetting Cluster Witness Server SSH configuration.
```

```
Verifying if the Cluster Witness has been configured on this Management Server...
```

```
Verifying if the Cluster Witness has been enabled on this Management Server...
```

```
VPN Reset between the Management Server and the Cluster Witness Server is now complete.
```

```
The log summary for configuration automation has been captured in
/var/log/Vplex/cli/Vplexconfig.log
```

```
The task summary and the commands executed for each automation task has been captured in
/var/log/Vplex/cli/Vplexcommands.txt
```

```
The output for configuration automation has been captured in
/var/log/Vplex/cli/capture/Vplexconfiguration-session.txt
```

From the second cluster:

```
Vplexcli:/> configuration cw-vpn-reset
```

```
This command will reset the vpn connectivity from this VPLEX instance to the Cluster Witness Server.
```

```
Do you want to continue? (Y/N): Y
```

To disable the VPN connectivity to the Cluster Witness Server please enter RESET (case sensitive):**RESET**  
Verifying if there is a VPN connection between the Management Server and the Cluster Witness Server...  
Verifying if the Cluster Witness has been configured on this Management Server...  
Verifying if the Cluster Witness has been enabled on this Management Server...  
  
Successfully removed the connection name and updated the Cluster Witness Server ipsec.conf file  
Successfully transferred the ipsec configuration file to the Cluster Witness Server and restarted the IPSec process  
  
Successfully removed the certificate files from the Cluster Witness Server  
Successfully removed the cluster witness connection name from the Management Server ipsec.conf file  
Successfully restarted the ipsec process on the Management Server  
  
Resetting Cluster Witness Server SSH configuration.  
Verifying if the Cluster Witness has been configured on this Management Server...  
Verifying if the Cluster Witness has been enabled on this Management Server...  
VPN Reset between the Management Server and the Cluster Witness Server is now complete.

The log summary for configuration automation has been captured in  
/var/log/VPLEX/cli/VPLEXconfig.log

The task summary and the commands executed for each automation task has been captured in  
/var/log/VPLEX/cli/VPLEXcommands.txt

The output for configuration automation has been captured in  
/var/log/VPLEX/cli/capture/VPLEXconfiguration-session.txt

- See also**
- ◆ [cluster-witness configure on page 96](#)
  - ◆ [cluster-witness enable on page 101](#)

## configuration enable-front-end-ports

After the meta-volume is created, continues the EZ-Setup wizard.

**Contexts** All contexts.

**Syntax** `configuration enable-front-end-ports`

**Description** Used to complete the initial system configuration using the EZ-Setup Wizard. After the meta-volume has been configured on the cluster, use this command to resume setup and enable the front-end ports on the local cluster.

Prerequisite: The cluster must be configured with a meta-volume and a meta-volume backup schedule.

**Example** `Vplexcli:/> configuration enable-front-end-ports`

**See also**

- ◆ [configuration continue-system-setup on page 114](#)
- ◆ [configuration metadata-backup on page 128](#)
- ◆ [configuration system-setup on page 135](#)
- ◆ [meta-volume create on page 306](#)

## configuration event-notices-reports config

Configure call-home and SYR settings after the initial configuration of VPLEX.

**Contexts** All contexts.

**Syntax** `configuration event-notices-reports-config`

**Description** This command runs an interview script that prompts for values to configure event notification (call-home) and reporting (SYR).

If both call-home and SYR are already configured, the current configuration information is displayed.

If either call-home or SYR is not configured, interview questions to configure the service that is not configured are displayed.

**Note:** This command does not modify an existing configuration. Use the “[configuration event-notices-reports reset](#)” command to reset (delete) an existing event notification and reporting configuration. Then use this command to configure new settings.

### Before you begin

You will need the following information to complete configuration of call-home and SYR reporting:

- ◆ IP address of the primary SMTP server used to forward reports to EMC. EMC recommends using your ESRS gateway as the primary connection address.
- ◆ (Optional) One or more IP address of secondary SMTP server(s) used to forward reports to EMC if the primary server fails. This address must be different than the address for the primary SMTP server.
- ◆ (Optional) One or more e-mail addresses of personnel who should receive e-mail notifications when events occur.

**Example** In the following example, neither call-home nor SYR is configured. The interview script prompts for information to configure both services.

**Note:** All IP addresses in the following example are for illustration only. Enter the IP addresses applicable to your configuration.

```
VPlexcli: /> configuration event-notices-reports-config
```

```
Configuring EMC Notification Options
```

```
By configuring EMC Event Notifications, EMC will be able to better serve you.
```

```
Would you like to configure this cluster to send event notifications to EMC? (yes/no)
[yes]yes
```

```
By sending EMC System Reports, EMC can proactively communicate known configuration risks to you. Any newly discovered information that can optimize or reduce risks to your system can also be communicated.
```

```
Would you like to configure this cluster to send system reports to EMC? (yes/no) [yes]yes
```

```
Configure Event and System Report Notifications
```

Events and/or system reports will be called home through an SMTP server. You will need to provide the SMTP IP v4 address for the primary connection. EMC recommends using your ESRS gateway as the primary connection address.

Enter the SMTP IP v4 address for the primary connection []**1.1.1.1**

Failover connection(s) will be attempted, in order, when the primary and any previous failover attempts have failed. This connection must go to EMC via a different SMTP server than the primary.

Would you like to configure a failover connection? (yes/no) [yes]**yes**

Enter the SMTP IP v4 address for this failover connection []**2.2.2.2**

Would you like to configure additional failover connections? (yes/no) [yes]**no**

Select Your Event Notification Options

By configuring your Event Notifications, people in your organization will receive email notifications when events occur. Events and/or system reports will be called home through an SMTP server. EMC recommends that you distribute your connections over multiple SMTP servers for better availability. Note that these SMTP v4 IP addresses can be different from the addresses used for EMC Event Notifications.

Would you also like one or more people in your organization to receive the events and/or system reports? (yes/no) [yes]**yes**

Notification Options

By configuring your Event Notifications, people in your organization will receive email notifications when events occur. Events and/or system reports will be called home through an SMTP server. EMC recommends that you distribute your connections over multiple SMTP servers for better availability. Note that these SMTP v4 IP addresses can be different from the addresses used for EMC Event Notifications.

1. On Success or Failure - Sends an email regardless of whether the email notification to EMC succeeded.
2. On Failure - Sends an email each time an attempt to notify EMC has failed.
3. On All Failure - Sends an email only if all attempts to notify EMC have failed.
4. On Success - Sends an email each time EMC is successfully sent an email notification.

Select your notification type (1 - 4) [1]**2**

Enter the recipient's email address []**test@emc.com**

Please re-enter the email address **test@emc.com**

Enter the SMTP IP v4 address for the connection []**3.3.3.3**

Would you like to notify additional members of your organization? (yes/no) [no]**no**

Configure EMC System Reports

You chose to send system reports back to EMC. Please note: All times are UTC and are not based on the local time.

VPLEX is configured to send a system report to EMC every Monday at 09:11.

Would you like to change the time the system report is sent? [no] **yes**

What day of the week should the system report be collected? (Sun=0..Sat=6) **5**

What hour of the day should the system report be collected? (0..23) **5**

What minute of the hour should the system report be collected? (0..59) **0**

VPLEX is configured to send a system report to EMC every Friday at 05:00.

Would you like to change the time the system report is sent? [no] **no**

Notification Information

Events and system reports will be sent by the cluster.

The Primary connection will go through SMTP server 1.1.1.1.

Failover connections will go through these SMTP servers:

2.2.2.2

Your organization will receive event notifications as follows:

aneuburg@emc.com will receive On Failure over 3.3.3.3

The system report will be sent home weekly on Friday at 05:00.

Would you like to run the setup process now? (yes/no) **yes**

Setting up environment for Call Home Connection Configuration...

Notification configuration started.

Verifying that all the pre-conditions for configuring the notification connection(s) are satisfied.

Verifying that the Call Home configuration file is present.

Verification of the pre-conditions for the run complete.

Configuring notification connection records.

Stopping EMC VPLEX connectemc

Stopped.

Starting EMC VPLEX connectemc

Started.

Stopping EMC VPLEXEmaAdaptor process

Stopped.

Starting EMC VPLEXEmaAdaptor process

Started.

SYR data collection job scheduled

Configuring notification connection records is completed successfully.

The call-home has been successfully configured

Vplexcli:/>

**Example** In the following example, both services have already been configured:

Vplexcli:/> **configuration event-notices-reports-config**

Call Home and System Report are already configured.

Notification Information

Events and system reports will be sent by the cluster.

The Primary connection will go through SMTP server 1.1.1.1.

Failover connections will go through these SMTP servers:

2.2.2.2

Your organization will receive event notifications as follows:

test@emc.com will receive On Failure over 3.3.3.3

The system report will be sent home weekly on Friday at 05:00.

**See also** ♦ [configuration event-notices-reports reset on page 124](#)

♦ [notifications call-home test on page 365](#)

♦ [scheduleSYR list on page 414](#)

♦ [Enable/disable call-home notifications on page 445](#)

## configuration event-notices-reports reset

Resets the current event notification and reporting configuration.

**Contexts** All contexts.

**Syntax** `configuration event-notices-reports-reset`

**Description** This command:

- ◆ Deletes the current event notification and reporting configuration data
- ◆ Disables event notification: call-home
- ◆ Removes the SYR schedule entry

**Example** In the following example:

- ◆ The `scheduleSYR list` and `ls /notifications/call-home/` commands display the current call-home and SYR reporting configuration.
- ◆ The `configuration event-notices-reports-reset` command removes the call-home and SYR reporting configuration.
- ◆ The `scheduleSYR list` and `ls /notifications/call-home/` commands confirm the change.

```
VPllexcli:/> scheduleSYR list
SYR data collection job is currently scheduled at:
Day of Week: 0 (Sunday=0, Monday=1,...Saturday=6)
Hours: 0
Minutes: 0
```

```
VPllexcli:/> ls /notifications/call-home/
```

```
/notifications/call-home:
```

```
Attributes:
```

```
Name      Value
-----  -
enabled   true
```

```
Contexts:
```

```
snmp-traps
```

```
VPllexcli:/> configuration event-notices-reports-reset
```

```
Call home is configured, Do you want to reset current configuration? (Y/N): y
```

```
Resetting all the tasks...
```

```
Cleaning up CallHomeAutomationTask...
```

```
Clearing all callhome records...
```

```
Restarting the call home services...
```

```
Stopping EMC Vplex connectemc           Stopped.
Starting EMC Vplex connectemc           Started.
Stopping EMC VplexEmaAdaptor process    Stopped.
Starting EMC VplexEmaAdaptor process    Started.
Disabling callhome...
```

```
Unsetting call home environment...
```

```
Resetting SYR schedule...
```

Reset of CallHomeAutomationTask is complete...

```
Vplexcli:/> scheduleSYR list
```

There is currently NO SYR data collection job scheduled.

```
Vplexcli:/> ls /notifications/call-home/
```

/notifications/call-home:

Attributes:

| Name    | Value |
|---------|-------|
| -----   | ----- |
| enabled | false |

Contexts:

snmp-traps

- See also**
- ◆ [configuration event-notices-reports config on page 121](#)
  - ◆ [notifications call-home test on page 365](#)
  - ◆ [scheduleSYR list on page 414](#)
  - ◆ [Enable/disable call-home notifications on page 445](#)

## configuration get-product-type

Displays the VPLEX product type (Local, Metro, or Geo).

**Contexts** All contexts.

**Syntax** `configuration get-product-type`

**Description** Displays whether the system is a Local, Metro, or Geo configuration.

**Example**  
VPllexcli:/> **configuration get-product-type**  
The cluster is currently configured as a VPLEX Metro

**See also**

- ◆ [cluster status on page 89](#)
- ◆ [cluster summary on page 93](#)
- ◆ [version on page 501](#)

---

## configuration join-clusters

Validates Fibre channel WAN connectivity and joins the two clusters.

**Contexts** All contexts.

**Syntax** `configuration join-clusters`

**Description** Completes the system configuration of a VPLEX Metro or Geo using the EZ-Setup Wizard.

**Example** `Vplexcli: /> configuration join-clusters`

- See also**
- ◆ [configuration connect-remote-directors on page 113](#)
  - ◆ [configuration continue-system-setup on page 114](#)
  - ◆ [configuration system-setup on page 135](#)
  - ◆ [verify fibre-channel-switches on page 500](#)

## configuration metadata-backup

Configures and schedules the daily backup of VPLEX metadata.

**Contexts** All contexts.

**Syntax** `configuration metadata-backup`

**Description** Selects the volumes to use as backup volumes and creates the initial backup of both volumes.

Runs an interview script that prompts for values to configure and schedule the daily backups of VPLEX metadata.

- ◆ Selects the volumes on which to create the backup
- ◆ Updates the VPlex configuration .xml file (VPlexconfig.xml)
- ◆ Creates an initial backup on both selected volumes
- ◆ Creates two backup volumes named:
  - *volume-1\_backup\_timestamp*
  - *volume-2\_backup\_timestamp*
- ◆ Schedules a backup at a time selected by the user

Enter two or more storage volumes, separated by commas.



### CAUTION

**Specify two or more storage volumes. Storage volumes must be:**  
**- unclaimed**  
**- on different arrays**

**Example** Configure the VPLEX metadata backup schedule:

```
VPlexcli:/clusters/cluster-1/system-volumes> configuration metadata-backup
```

Configuring Meta-data Backups

To configure meta-data backups you will need to select two volumes (78G or greater), preferably on two different arrays. Backups will occur automatically each day, at a time you specify.

Available Volumes for Meta-data Backup

| Name                                     | Capacity | Vendor | IO Status | Type        |
|------------------------------------------|----------|--------|-----------|-------------|
| VPD83T3:60000970000192601714533036464236 | 80.1G    | EMC    | alive     | traditional |
| VPD83T3:60000970000192601714533036464237 | 80.1G    | EMC    | alive     | traditional |

Please select volumes for meta-data backup, preferably from two different arrays

(volume1,volume2): **VPD83T3:60000970000192601714533036464236,VPD83T3:60000970000192601714533036464237**

What hour of the day should the meta-data be backed up? (0..23): **11**

What minute of the hour should the meta-data be backed up? (0..59): **25**

VPLEX is configured to back up meta-data every day at 11:25 (UTC).

Would you like to change the time the meta-data is backed up? [no]: **no**

Review and Finish

Review the configuration information below. If the values are correct, enter

yes (or simply accept the default and press Enter) to start the setup process. If the values are not correct, enter no to go back and make changes or to exit the setup.

Meta-data Backups

Meta-data will be backed up every day at 11:25.

The following volumes will be used for the backup :

VPD83T3:60000970000192601714533036464236,

VPD83T3:60000970000192601714533036464237

Would you like to run the setup process now? [yes]:

**Example** Modify the existing daily backup of VPLEX metadata:

```

Vplexcli:/clusters/cluster-1/system-volumes> configuration metadata-backup
A back up of the meta-data is already scheduled to occur everyday at 11:25 (UTC). Do you want
change the existing schedule? (Y/N): y
Configuring Meta-data Backups
To configure meta-data backups you will need to select two volumes (78G or greater), preferably
on two different arrays. Backups will occur automatically each day, at a time you specify.

VPLEX is currently configured to backup metadata on the following volumes :
VPD83T3:60000970000192601714533036464236,
VPD83T3:60000970000192601714533036464237

Would you like to change the volumes on which to backup the metadata? [no]:
VPLEX is configured to back up meta-data every day at 11:25 (UTC).
Would you like to change the time the meta-data is backed up? [no]: yes
What hour of the day should the meta-data be backed up? (0..23): 11
What minute of the hour should the meta-data be backed up? (0..59): 00
VPLEX is configured to back up meta-data every day at 11:00 (UTC).
Review and Finish
  Review the configuration information below. If the values are correct, enter
  yes (or simply accept the default and press Enter) to start the setup process. If the values
  are not correct, enter no to go back and make changes or to exit the setup.
    Meta-data Backups
      Meta-data will be backed up every day at 11:20.
      The following volumes will be used for the backup :
VPD83T3:60000970000192601714533036464236,
VPD83T3:60000970000192601714533036464237
    Would you like to run the setup process now? [yes]: yes

```

Use the `ls /clusters/cluster-2/system-volumes/` command to display the backup meta-volumes:

```

Vplexcli:/> ls /clusters/cluster-2/system-volumes/

/clusters/cluster-2/system-volumes:
Detroit_LOGGING_VOL_vol Detroit_METAVolume1 Detroit_METAVolume1_backup_2010Dec23_052818
Detroit_METAVolume1_backup_2011Jan16_211344

```

- See also**
- ◆ [configuration show-meta-volume-candidates on page 132](#)
  - ◆ [configuration system-setup on page 135](#)

## configuration register-product

Registers the VPLEX product with EMC.

**Contexts** All contexts.

**Syntax** `configuration register-product`

**Description** Use this command during installation:

- ◆ After configuring the external IP address and host name
- ◆ Before using the “`configuration system-setup`” command.

Runs the product registration wizard. Prompts for the following information:

- ◆ Company contact name
- ◆ E-mail and phone number
- ◆ Mailing address.

The command uses the responses to create a file for product registration. A prompt is then displayed asking how the registration should be sent to EMC. Two methods are available:

- ◆ Attach the registration file to an e-mail and send it to EMC.
- ◆ Send the registration file to EMC through an SMTP server.

If this option is selected, a prompt for an SMTP server IP address is displayed.

### Example

```
VPlxcli:/> configuration register-product
```

```
Welcome to the VPLEX Product Registration Assistant. To register your
VPLEX product, please provide the information prompted for below.
This information will be sent to EMC via email or, this information will be captured in a file
that you can attach to an email to send to EMC.
```

```
Attempting to determine the VPLEX Product Serial Number from the system.
```

```
Company Site ID Number (Optional) : 12345
Company Name : EMC Test
Contact First Name : Customer
Contact Last Name : One
Contact Email Address : customer.one@company.com
Contact Phone Number : 888-555-1212
Company Street Address : 176 Main Street
Company City : Boston
Company State or Province : MA
Company Zip or Postal Code : 01748
Company Country : USA
```

```
Which method will be used to Connect Home. Enter the number associated with your selection.
```

```
1: ESRS 2: Email Home 3: Do Not Connect Home
```

```
Connect Home using : 3
```

```
Which method will be used for Remote Support. Enter the number associated with your selection.
```

```
1: ESRS 2: WebEx
```

```
Remote Support using : 2
```

```
Please review your product registration information below.
```

```
VPLEX Product Serial Number : DEV12345678
VPLEX software version : 0.17.25.0.0
Company Site ID # : 12345
Company Name : EMC Test
First Name : Customer
```

Last Name : One  
Business Email Address : customer.one@company.com  
Business Phone Number : 888-555-1212  
Business Address : 176 Main Street  
City : Boston  
State/Province : MA  
Zip/Postal Code : 01748  
Country : USA  
Connect Home Using : Do not Connect Home  
Remote Support Using : WebEx

Would you like to continue with these values? (Y/N): **y**

To complete the registration process, this information must be sent to EMC. We can send this product registration information to EMC for you using an SMTP server of your choice. Would you like to send this now?

(Y/N): **n**

To complete the registration process, this information must be sent to EMC. Please attach the file located here:  
/var/log/Vplex/cli/productRegistration.txt  
to an email and send it to b2b\_product\_registrations@emc.com as soon as possible to complete your VPLEX registration.

- See also**
- ◆ [configuration continue-system-setup on page 114](#)
  - ◆ [configuration system-setup on page 135](#)

## configuration show-meta-volume-candidates

Display the volumes which meet the criteria for a VPLEX meta-volume.

**Contexts** All contexts.

**Syntax** `configuration show-meta-volume-candidates`

**Description** Candidate volumes are:

- ◆ Unclaimed
- ◆ At least 78 Gb capacity



### CAUTION

**If the meta-volume is configured on a CLARiiON® array, it must not be placed on the vault drives of the CLARiiON.**

EMC recommends the following for meta-volumes:

- ◆ Read caching should be enabled
- ◆ A hot spare meta-volume be pre-configured in case of a catastrophic failure of the active meta-volume.

Performance is not critical for meta-volumes. The minimum performance allowed is 40 MB/s and 100 4 K IOP/second. The physical spindles for meta-volumes should be isolated from application workloads.

Availability IS critical for meta-volumes. Best practice is to mirror the meta-volume across two or more back-end arrays. Choose the arrays used to mirror the meta-volume such that they are not required to migrate at the same time.

### Example

```
Vplexcli:/> configuration show-meta-volume-candidates
```

| Name                                     | Capacity | Vendor | IO Status | Type        | Array Name              |
|------------------------------------------|----------|--------|-----------|-------------|-------------------------|
| VPD83T3:60060480000190100547533030364539 | 187G     | EMC    | alive     | traditional | EMC-SYMMETRIX-190100547 |
| VPD83T3:60000970000192601707533031333132 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333133 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333134 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333135 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333136 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333137 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:60000970000192601707533031333138 | 98.5G    | EMC    | alive     | traditional | EMC-SYMMETRIX-192601707 |
| VPD83T3:6006016049e02100442c66c8890ee011 | 80G      | DGC    | alive     | traditional |                         |
| EMC-CLARiiON-FNM00083800068              |          |        |           |             |                         |

The log summary for configuration automation has been captured in  
/var/log/Vplex/cli/Vplexconfig.log

The task summary and the commands executed for each automation task has been captured in  
/var/log/Vplex/cli/Vplexcommands.txt

The output for configuration automation has been captured in  
/var/log/Vplex/cli/capture/Vplexconfiguration-session.txt  
Vplexcli:/>

- See also**
- ◆ [configuration metadata-backup on page 128](#)
  - ◆ [configuration system-setup on page 135](#)
  - ◆ [meta-volume create on page 306](#)

## configuration sync-time

Synchronizes the time of the local management server with a remote management server.

**Syntax** `configuration sync-time`

**Description** In a VPLEX Metro or Geo configuration, synchronizes the time between the local management server and a remote management server using NTP.



### **CAUTION**

**May cause the CLI or SSH session to disconnect. If this occurs, re-log in and continue system set-up where you left off.**

Use this command only on the second cluster to be configured in a VPLEX Metro or Geo configuration.

Use this command before any set-up has been performed on the second cluster of a VPLEX Metro or Geo configuration.

Use this command during initial system configuration before using the `configuration system-setup` command.

Use this command only on the second cluster to be configured in a VPLEX Metro or Geo configuration.

This command synchronizes Cluster 2 with Cluster 1 even if the time was correct on Cluster 2 but incorrect on Cluster 1.

The command prompts for the public IP address of the first cluster's management server.

**Example** `Vplexcli:/> configuration sync-time`

Please enter the IP address of the remote management server with which you want to synchronize the system clock.

**See also** ♦ [configuration system-setup on page 135](#)

## configuration system-reset

Resets the cluster to the manufacturing state.

**Contexts** All contexts.

**Syntax** `configuration system-reset`

**Description** Reset the cluster to the manufacturing state. Usage varies by release:

- ◆ For Release 4.1, use this command to cancel all the configuration completed using the EZ-Setup wizard and return the VPLEX to its factory default settings. Any values specified during the configuration session become the defaults displayed when EZ-Setup is re-run.
- ◆ For Release 4.2, use the **exit** command to exit the EZ-Setup wizard.

Run the **configuration system-reset** command only on a new cluster.

Do not run this command on a configured cluster.

No meta-volume may exist for this command to execute.

**Example** `VPllexcli:/> configuration system-reset`

**See also** ◆ [configuration system-setup on page 135](#)

## configuration system-setup

Starts the EZ-Setup Wizard automated configuration tool.

**Contexts** All contexts.

**Syntax** `configuration system-setup`

**Arguments** **Required arguments**  
**[-m | --setup-multiplex-specific]** - Use this argument only on the first cluster added to a VPLEX Metro or Geo configuration.

Configures the VPN and establishes a secure connection between the clusters

Before using this argument, use the **configuration system-setup** command on the second cluster of a VPLEX Metro or Geo configuration. Refer to the *EMC VPLEX Installation and Setup Guide*.

**Description** Use the **exit** command any time during the session to exit EZ-Setup.  
 Use the **configuration system-setup** command to resume EZ-Setup. The process restarts from the first step. Any values from the previous session appear as default values.



### **IMPORTANT**

**There must be no meta-volume configured, and no storage exposed to hosts in order to run this command.**

EZ-Setup Wizard automates the following steps:

- ◆ Connects to the local directors
- ◆ Sets the cluster ID
- ◆ Sets the Cluster IP Seed (see [About cluster IP seed and cluster ID on page 434](#))
- ◆ Sets the director count
- ◆ Commissions the directors
- ◆ Configures NTP
- ◆ Configures Fibre Channel switch
- ◆ Enables COM and back-end ports
- ◆ Creates Web server and VPN certificate
- ◆ Configures VPN
- ◆ Configures/enables
- ◆ Configures/enables System Reporting (SYR)

### **General configuration procedure**

During initial installation, the configuration commands are generally used in the following order:

- ◆ Use the “[configuration system-setup](#)” command to start the EZ-Setup Wizard. Configure the VPLEX based on information typed in response to a series of questions.

Use the **exit** command to exit the EZ-Setup Wizard.

Use the “[configuration system-setup](#)” command to restart the wizard from the first step. Values entered in the previous session are displayed as the defaults.

- ◆ Zone the back-end storage to the port WWNs of the VPLEX back-end ports.
- ◆ Use the “[configuration continue-system-setup](#)” command to complete the initial configuration.
- ◆ Use the “[configuration show-meta-volume-candidates](#)” command to display storage volumes that are candidates for the meta-volume (unclaimed, no extents, and at least 78 GB capacity).
- ◆ Create the meta-volume.
- ◆ Use the “[configuration enable-front-end-ports](#)” command to enable the front-end ports on the local cluster.

For VPLEX Metro or Geo configurations:

- ◆ Use the “[configuration continue-system-setup](#)” command to complete the configuration for the local cluster.
- ◆ Enable and zone the WAN ports.
- ◆ Use the “[configuration connect-remote-directors](#)” command to connect the local cluster to the directors on the remote cluster, and the remote cluster to the directors on the local cluster.

After configuration is complete, use the following commands to make modifications:

- ◆ Use the “[configuration configure-auth-service](#)” to configure the authentication service selected by the user.
- ◆ Use the “[configuration cw-vpn-configure](#)” and “[configuration cw-vpn-reset](#)” commands to manage VPLEX Witness.
- ◆ Use the “[configuration event-notices-reports config](#)” and “[configuration event-notices-reports reset](#)” commands to import/apply/remove customized call-home event files.

**Example** `Vplexcli:/> configuration system-setup`

- See also**
- ◆ [About cluster IP seed and cluster ID on page 434](#)
  - ◆ [configuration continue-system-setup on page 114](#)
  - ◆ [configuration system-reset on page 134](#)

## connect

Connects to a director.

**Contexts** All contexts.

**Syntax**

```
connect
[-o|--host] [host name|IP address]
--logport port number
--secondary-host [host name|IP address]
--secondary-logport secondary port number
[-n|--name] name
[-t|--type] system type
[-p|--password] password
[-c|--connection-file] filename
[-s|--save-authentication]
--no-prompt
```

**Arguments** **Required arguments**

None.

### Optional arguments

**[-o|--host] {host-name | IP address}** - \* Hostname or IP address of the director to which to connect.

Default: localhost.

**--logport port-number** - For use by EMC personnel only. A firmware log event port. Applicable only to test versions of the firmware.

**--secondary-host {host-name | IP address}** - Hostname or IP address of the redundant interface on director to which to connect.

**--secondary-logport secondary-port-number** - For use by EMC personnel only. A firmware log event port. Applicable only to test versions of the firmware.

**--name name** - \* Name to assign to the director after the connection.

Default: Hostname or IP address.

**[-t|--type] type** - For use by EMC personnel only. VPLEX can communicate with its firmware through two styles of interfaces. Tools such as the VPLEX simulator use the 'legacy' type interface.

**[-p|--password] password** - Set the password for the connection.

**[-c|--connection-file] filename** - Load a list of connections from file named 'connections' at:

service@ManagementServer:/var/log/VPlex/cli on the VPLEX management server.

**[-s|--save-authentication]** - Save the authentication information used to connect to the VPLEX without asking for confirmation.

**--no-prompt** - Do not prompt for a password when authentication to the VPLEX is required.

\* - argument is positional.

**Description** Use the **connect** command to:

- ◆ Re-establish connectivity if connectivity is lost to one or more directors.
- ◆ Manually re-connect after a power outage when the management server boots up before the directors are available.

After a power outage, if the management server boots before the directors are available, the CLI is operational before the directors can respond to the automatic connect. Use the **connect -c *connection-filename*** command to manually re-connect the directors.

During normal system setup, connections to directors are established and stored in a file: `/var/log/VPLEX/cli/connections`.

Use the **connect -c *filename*** command if the entry for the director exists in the connections file.

**Note:** If the disconnect command is issued for a director, the entry in the connections file for that director is removed from the connections file.

When a director is connected, the context tree expands with new contexts representing the director, including:

- ◆ A new director context below `/engines/engine/directors` representing storage and containing the director's properties.
- ◆ If this is the first connection to a director at that cluster, a new cluster context below `/clusters`.
- ◆ If this is the first connection to a director belonging to that engine, a new engine context below `/engines`.

Use the **connect -name *name*** command to name the new context below `/engines/engine/directors`. If is omitted, the hostname or IP address is used.

**Note:** *name* is required if two host addresses are specified.

If the **--connection-file** argument is used, the specified file must list at least one host address to connect to on each line with the format:

```
host | [secondary host][,name]
```

Sample connections file:

```
service@ManagementServer:/var/log/VPLEX/cli> tail connections
128.221.252.67:5988|128.221.253.67:5988,Cluster_2_Dir_1A
128.221.252.68:5988|128.221.253.68:5988,Cluster_2_Dir_1B
128.221.252.69:5988|128.221.253.69:5988,Cluster_2_Dir_2A
128.221.252.70:5988|128.221.253.70:5988,Cluster_2_Dir_2B
128.221.252.35:5988|128.221.253.35:5988,Cluster_1_Dir1A
128.221.252.36:5988|128.221.253.36:5988,Cluster_1_Dir1B
128.221.252.36:5988,128.221.252.36
128.221.252.35:5988,128.221.252.35
```

**Example** `Vplexcli:> connect --host 128.221.252.67 --secondary-host 128.221.253.67 --name director-2-1-A`

- See also**
- ◆ [configuration connect-local-directors on page 112](#)
  - ◆ [disconnect on page 206](#)

## connectivity director

Displays connections from the specified director through data (non-management) ports.

**Contexts** All contexts.

**Syntax** `connectivity director`  
*director*  
 [-d|--storage-volumes]  
 [-i|--initiators]  
 [-n|--directors]  
 [-f|--file] *filename*  
 [-s|sort-by] [*name|wwn|port*]

**Arguments** **Required arguments**  
*director* - Director to discover.

### Optional arguments

**[-d|--storage-volumes]** - Display connectivity from the specified director to storage volumes.

**[-i|--initiators]** - Display connectivity from the specified director to initiators.

**[-n|--directors]** - Display connectivity from the specified director to other directors.

**[-f|--file] filename** - Save the output in the specified file.

Default: /var/log/VPLEX/cli

**[-s|--sort-by] {name|wwn|port}** - Sort output by one of the following:

**name** - Sort output by storage volume name.

**wwn** - Sort output by WorldWide name.

**port** - Sort output by port.

**Description** Prints a table of discovered storage volumes, initiators and directors. Lists the ports on which it discovered each storage volume, initiator and director.

**Example** Display connectivity from director Cluster\_1\_DirA to initiators:

```
Vplexcli:/> connectivity director Cluster_1_Dir1A/ --initiators
Initiators discovered
Node WWN                Port WWN                Ports
-----
0x20000000c992c3d5    0x10000000c992c3d5    A0-FC00,A1-FC00
0x20000000c97b1f3d    0x10000000c97b1f3d    A0-FC00,A1-FC00
0x20000000c97b1f3c    0x10000000c97b1f3c    A0-FC02,A1-FC02
0x20000000c97b0c3f    0x10000000c97b0c3f    A0-FC00,A1-FC00
0x20000000c97b0c3e    0x10000000c97b0c3e    A0-FC02,A1-FC02
0x20000000c990d6fd    0x10000000c990d6fd    A0-FC02,A1-FC02
0x20000000c992c219    0x10000000c992c219    A0-FC02,A1-FC02
0x20000000c992c841    0x10000000c992c841    A0-FC00,A1-FC00
```

**Example** Display connectivity from director Cluster\_1\_DirA to directors:

```
Vplexcli:/> connectivity director Cluster_1_Dir1A/ --directors
Device VPD83T3:6000970000192601378533030303530 is a default LUN_0.
Device VPD83T3:6000970000192601723533030303530 is a default LUN_0.
Device VPD83T3:6000970000192601852533030303530 is a default LUN_0.
Directors discovered by Cluster_1_Dir1A, UUID 0x000000003ca00147:
Director UUID          Protocol Address          Ports
```

```

-----
0x000000003ca000e6 COMSCSI 0x500014424000e643 A4-FC03,A4-FC02
0x000000003cb001cb COMSCSI 0x500014425001cb43 A4-FC02,A4-FC03
0x000000003ca001cb COMSCSI 0x500014424001cb43 A4-FC02,A4-FC03
0x000000003cb00147 COMSCSI 0x5000144250014741 A4-FC01
                                COMSCSI 0x5000144250014740 A4-FC00
0x000000003cb000e6 COMSCSI 0x500014425000e643 A4-FC02,A4-FC03

```

**Example** Display connectivity from director Cluster\_1\_DirA and sort output by port:

```

VPLEXcli:/> connectivity director Cluster_1_Dir1A/ --sort-by port
Device VPD83T3:60000970000192601378533030303530 is a default LUN_0.
Device VPD83T3:60000970000192601723533030303530 is a default LUN_0.
Device VPD83T3:60000970000192601852533030303530 is a default LUN_0.
StorageVolumes discovered - sorted by: port
Port      WWN                StorageVolume Name                LUN
-----
A2-FC00  0x50000972081aed5  VPD83T3:60000970000192601723533030313530  0x0001000000000000
                                VPD83T3:60000970000192601723533030313534  0x0002000000000000
                                0x500601603ce03506  VPD83T3:6006016021d0250026b925ff60b5de11  0x0008000000000000
                                VPD83T3:6006016021d0250027b925ff60b5de11  0x0001000000000000
.
.
.
A2-FC02  0x500601613ce03506  VPD83T3:6006016021d0250026b925ff60b5de11  0x0008000000000000
                                VPD83T3:6006016021d0250027b925ff60b5de11  0x0001000000000000
.
.
.

```

**Example** Display connectivity from director Cluster\_1\_DirA and sort output by storage volume name:

```

VPLEXcli:/> connectivity director Cluster_1_Dir1A/ --sort-by name
Device VPD83T3:60000970000192601378533030303530 is a default LUN_0.
Device VPD83T3:60000970000192601723533030303530 is a default LUN_0.
Device VPD83T3:60000970000192601852533030303530 is a default LUN_0.
StorageVolumes discovered - sorted by: name
StorageVolume Name                WWN                LUN                Ports
-----
VPD83T3:60000970000192601723533030313530  0x50000972081aed5  0x0001000000000000  A2-FC00,A3-FC00
VPD83T3:60000970000192601723533030313534  0x50000972081aed5  0x0002000000000000  A2-FC00,A3-FC00
VPD83T3:6006016021d0250026b925ff60b5de11  0x500601603ce03506  0x0008000000000000  A2-FC00,A3-FC00
                                0x500601613ce03506  0x0008000000000000  A2-FC02,A3-FC02
.
.
.

```

**See also**

- ◆ [connectivity show on page 141](#)
- ◆ [connectivity validate-be on page 142](#)

## connectivity show

Displays the communication endpoints that can see each other.

**Contexts** All contexts.

**Syntax** `connectivity show`  
`[-p|--protocol] {fc|ib|tcp|udt}`  
`[-e|--endpoints] port,port...`

**Arguments** **Required arguments**  
 None.

### Optional arguments

`[-p|--protocol] {fc|ib|tcp|udt}` - Display endpoints with only the specified protocol. Arguments are case-sensitive, and include:

**fc** - Fibre Channel.

**ib** - InfiniBand. Not supported in the current release. Use the “[connectivity director](#)” command to display IB protocol connectivity.

**tcp** - Transmission Control Protocol.

**udt** - UDP-based Data Transfer Protocol.

`[-e|--endpoints] port,port...` - List of one or more ports for which to display endpoints. Entries must be separated by commas.

Default: Display endpoints for all ports.

**Description** Displays connectivity, but does not perform connectivity checks. Displays which ports can talk to each other.

**Example** Show all connectivity:

```
Vplexcli: /> connectivity show
/engines/engine-1-1/directors/Cluster_1_Dir1A/hardware/ports/A4-FC00 ->
/engines/engine-1-1/directors/Cluster_1_Dir1B/hardware/ports/B4-FC00
/engines/engine-1-1/directors/Cluster_1_Dir1A/hardware/ports/A4-FC01 ->
/engines/engine-1-1/directors/Cluster_1_Dir1B/hardware/ports/B4-FC01
.
.
.
```

**Example** Use the `--endpoints` argument followed by the <Tab> key to display all ports:

```
Vplexcli: /> connectivity show --endpoints
A0-FC00/  A0-FC01/  A0-FC02/  A0-FC03/  A1-FC00/  A1-FC01/  A1-FC02/
A1-FC03/  A2-FC00/  A2-FC01/  A2-FC02/  A2-FC03/  A3-FC00/  A3-FC01/
.
.
.
```

**Example** Show connectivity to two specified ports:

```
Vplexcli: /> connectivity show --endpoints A4-FC01/,B4-FC02/
/engines/engine-1-1/directors/Cluster_1_Dir1B/hardware/ports/B4-FC02 ->
/engines/engine-2-1/directors/Cluster_2_Dir_1B/hardware/ports/B4-FC02
/engines/engine-2-2/directors/Cluster_2_Dir_2B/hardware/ports/B4-FC02
.
.
.
```

**See also** ♦ [connectivity director on page 139](#)

## connectivity validate-be

Checks that the back-end connectivity is correctly configured.

**Contexts** All contexts.

**Syntax** `connectivity validate-be`  
`[-s|--summaryonly]`

**Arguments** **Required arguments**  
 None.

**Optional arguments**

`[-s|--summaryonly]` - Display only a summary. Default is to display details, followed by the summary.

**Description** This command checks the following:

- ◆ All directors see the same set of storage volumes.
- ◆ All directors have at least two paths to each storage-volume.
- ◆ The number of active paths from each director to a storage volume does not exceed 4.



**IMPORTANT**

**If the number of paths per storage volume per director exceeds 8 a warning event, but not a call home is generated. If the number of paths exceeds 16, an error event and a call-home notification are generated.**

On VPLEX Metro systems where RecoverPoint is deployed, run this command on both clusters.

If the “[connectivity director](#)” command is run for every director in the VPLEX prior to running this command, this command displays an analysis/summary of the back-end connectivity information.

**Example** Display validation for back-end connectivity on a healthy VPLEX Metro:

```
Vplexcli:/> connectivity validate-be
Summary
Cluster cluster-1
  0 storage-volumes which are dead or unreachable.
  0 storage-volumes which do not meet the high availability requirement for storage volume
paths*.
  0 storage-volumes which are not visible from all directors.
  0 storage-volumes which have more than supported (4) active paths from same director.
  *To meet the high availability requirement for storage volume paths each storage volume
must be accessible from each of the directors through 2 or more Vplex backend ports, and 2 or
more Array target ports, and there should be 2 or more ITLs.

Cluster cluster-2
  0 storage-volumes which are dead or unreachable.
  0 storage-volumes which do not meet the high availability requirement for storage volume
paths*.
  0 storage-volumes which are not visible from all directors.
  0 storage-volumes which have more than supported (4) active paths from same director.
  *To meet the high availability requirement for storage volume paths each storage volume
must be accessible from each of the directors through 2 or more Vplex backend ports, and 2 or
more Array target ports, and there should be 2 or more ITLs.
```

**Example** Storage summary when there are too many active paths from one director to one or more storage volumes:

```
Vplexcli: /> connectivity validate-be
```

```
Summary
```

```
Cluster cluster-1
```

```
.
.
.
```

```
Storage:
```

| Cluster Name | Total Storage Volumes | Unhealthy Storage Volumes | Total Virtual Volumes | Unhealthy Virtual Volumes | Total Dist Devs | Unhealthy Dist Devs | No Dual Paths | Not visible from All Dirs | With Unsupported # of Paths |
|--------------|-----------------------|---------------------------|-----------------------|---------------------------|-----------------|---------------------|---------------|---------------------------|-----------------------------|
| cluster-1    | 204                   | 0                         | 0                     | 0                         | 0               | 0                   | 0             | 0                         | 0                           |
| cluster-2    | 204                   | 0                         | 0                     | 0                         | 0               | 0                   | 204           | 0                         | 204                         |

WARNING: Storage volumes which have more than 1 active paths from the same director:

| Cluster   | Director       | Array                                | Storage volumes which have more than 1 active paths from the same director           |
|-----------|----------------|--------------------------------------|--------------------------------------------------------------------------------------|
| cluster-2 | director-2-1-A | (192601426, EMC~SYMMETRIX~192601426) | VPD83T3:60000970000192601426533030443034<br>VPD83T3:60000970000192601426533030443033 |

**Example** Display back-end connectivity in a VPLEX Geo with:

- ◆ VPLEX Witness deployed
- ◆ RecoverPoint delayed
- ◆ The number of active paths from some directors to storage volumes exceeds the limit of 4:

```
Vplexcli: /> connectivity validate-be
```

```
Summary
```

```
Cluster cluster-1
```

```
0 storage-volumes which are dead or unreachable.
```

```
0 storage-volumes which do not meet the high availability requirement for storage volume paths*.
```

```
0 storage-volumes which are not visible from all directors.
```

```
1556 storage-volumes which have more than 1 active paths from same director.
```

```
Director director-1-1-A
```

```
Storage array: ('APM00114103169', 'EMC~CLARiION~APM00114103169') has 3 storage-volumes which have more than 1 active paths from same director.
```

```
Storage array: ('190300487', 'EMC~SYMMETRIX~190300487') has 553 storage-volumes which have more than 1 active paths from same director.
```

```
Storage array: ('14403b', 'EMC~Invista~14403b') has 1000 storage-volumes which have more than 1 active paths from same director.
```

```
Director director-1-1-B
```

```
Storage array: ('APM00114103169', 'EMC~CLARiION~APM00114103169') has 3 storage-volumes which have more than 1 active paths from same director.
```

```
.
.
.
```

\*To meet the high availability requirement for storage volume paths each storage volume must be accessible from each of the directors through 2 or more Vplex backend ports, and 2 or more Array target ports, and there should be 2 or more ITLs.

```
Cluster cluster-2
```

```
0 storage-volumes which are dead or unreachable.
```

```
.
.
```

.  
Storage:

| Cluster Name | Total Storage Volumes | Unhealthy Storage Volumes | Total Virtual Volumes | Unhealthy Virtual Volumes | Total Dist Devs | Unhealthy Dist Devs | No Dual Paths | Not visible from All Dirs | With Excessive Paths |
|--------------|-----------------------|---------------------------|-----------------------|---------------------------|-----------------|---------------------|---------------|---------------------------|----------------------|
| cluster-1    | 1556                  | 0                         | 589                   | 0                         | 64              | 0                   | 0             | 0                         | 1556                 |
| cluster-2    | 1067                  | 0                         | 599                   | 0                         | 64              | 0                   | 0             | 0                         | 1067                 |

Consistency Groups:

| Cluster Name | Total Synchronous Groups | Unhealthy Synchronous Groups | Total Asynchronous Groups | Unhealthy Asynchronous Groups |
|--------------|--------------------------|------------------------------|---------------------------|-------------------------------|
| cluster-1    | 1                        | 0                            | 2                         | 0                             |
| cluster-2    | 1                        | 0                            | 2                         | 0                             |

IP WAN Connectivity:

| Cluster Name | Port Group   | Subnet         | Local Cluster Ips | Remote Cluster Ips | MTU  | Connectivity |
|--------------|--------------|----------------|-------------------|--------------------|------|--------------|
| cluster-1    | port-group-0 | cluster-1-SN00 | 192.168.11.251    | 192.168.11.252*    | 1500 | ok           |
| cluster-2    | port-group-0 | cluster-2-SN00 | 192.168.11.252    | 192.168.11.251*    | 1500 | ok           |

\* remote-ip-address is not specified in subnets context

Cluster Witness:

| Admin State | Private Ip    | Public Ip   | cluster-1 OP State | cluster-2 OP State | server OP State |
|-------------|---------------|-------------|--------------------|--------------------|-----------------|
| disabled    | 128.221.254.3 | 10.6.209.48 | -                  | -                  | -               |

RecoverPoint:

| Cluster Name | Total RP Clusters | Unhealthy RP Clusters | Total Replicated Virtual Volumes | Unhealthy Replicated Virtual Volumes | Total RP-enabled Consistency Groups | Mis-aligned RP-enabled Consistency Groups | Total Registered RP Initiators/Storage Views | Unhealthy RP Storage Views |
|--------------|-------------------|-----------------------|----------------------------------|--------------------------------------|-------------------------------------|-------------------------------------------|----------------------------------------------|----------------------------|
| cluster-1    | 0                 | 0                     | 0                                | 0                                    | 0                                   | 0                                         | 0/0                                          | 0                          |
| cluster-2    | -                 | -                     | -                                | -                                    | -                                   | -                                         | -/-                                          | -                          |

\*\*This command is only able to check the health of the local cluster(cluster-1)'s RecoverPoint configuration, therefore if this system is a VPLEX Metro or VPLEX Geo repeat this command on the remote cluster to get the health of the remote cluster's RecoverPoint configuration.

- See also**
- ◆ [connectivity director on page 139](#)
  - ◆ [connectivity show on page 141](#)
  - ◆ [connectivity validate-be on page 142](#)
  - ◆ [health-check on page 273](#)
  - ◆ [rp validate-configuration on page 405](#)
  - ◆ [validate-system-configuration on page 492](#)

## connectivity validate-wan-com

Verifies the expected IP and FC WAN COM connectivity.

**Contexts** All contexts.

**Syntax** `connectivity validate-wan-com`  
`[-e|--show-expected]`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-e|--show-expected]` - Display the expected connectivity map instead of comparing it to the actual connectivity.

The map is a list of every port involved in the WAN COM network and the ports to which it is expected to have connectivity.

**Description** This command assembles a list of expected WAN COM connectivity, compares it to the actual WAN COM connectivity and reports any discrepancies (i.e. missing or extra connections).

This command verifies IP or FC based WAN COM connectivity.

If no option is specified, displays a list of ports that are in error: either missing expected connectivity or have additional unexpected connectivity to other ports.

The expected connectivity is determined by collecting all ports with role wan-com and requiring that each port in a group at a cluster have connectivity to every other port in the same group at all other clusters.

When both FC and IP ports with role wan-com are present, the smaller subset is discarded and the protocol of the remaining ports is assumed as the correct protocol.

**Example** Display a map of expected WAN connectivity:

```
VPllexcli: /> connectivity validate-wan-com -e
Expected connectivity:
```

```
port-group-1
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG01 ->
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG01
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG01
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG01 ->
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG01
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG01
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG01 ->
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG01
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG01
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG01 ->
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG01
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG01

port-group-0
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG00 ->
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG00
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG00
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG00 ->
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG00
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG00
/engines/engine-2-1/directors/director-2-1-A/hardware/ports/A2-XG00 ->
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG00
```

```
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG00
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG00 ->
/engines/engine-1-1/directors/director-1-1-A/hardware/ports/A2-XG00
```

**Example** Display/verify WAN connectivity when ports are healthy, connectivity is correctly configured:

```
Vplexcli:/> connectivity validate-wan-com
connectivity: FULL
```

```
port-group-1 - OK - All expected connectivity is present.
port-group-0 - OK - All expected connectivity is present.
```

**Example** Display/verify WAN connectivity when there are errors communicating with directors or with the port-group configuration:

```
Vplexcli:/> connectivity validate-wan-com
connectivity: PARTIAL
```

```
port-group-1 - OK - All expected connectivity is present.
```

```
port-group-0 - ERROR - Connectivity errors were found for the following com ports:
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG00 ->
Missing connectivity to /engines/engine-2-2/directors/director-2-2-B/hardware/ports/B2-XG00
```

**Example** Display/verify WAN connectivity when no connectivity was detected between any ports:

```
Vplexcli:/> connectivity validate-wan-com
connectivity: NONE
```

```
port-group-1 - ERROR - Connectivity errors were found for the following com ports:
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG01 ->
Missing connectivity to /engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG01
/engines/engine-2-1/directors/director-2-1-B/hardware/ports/B2-XG01 ->
Missing connectivity to /engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG01
```

```
port-group-0 - ERROR - Connectivity errors were found for the following com ports:
/engines/engine-1-1/directors/director-1-1-B/hardware/ports/B2-XG00 ->
Missing connectivity to /engines/engine-2-2/directors/director-2-2-B/hardware/ports/B2-XG00
/engines/engine-1-2/directors/director-1-2-A/hardware/ports/A2-XG00 ->
Missing all expected connectivity.
```

**See also**

- ◆ [connectivity director on page 139](#)
- ◆ [connectivity show on page 141](#)
- ◆ [connectivity validate-be on page 142](#)

## connectivity window set

Sets values that control the operation of communications.



### CAUTION

**This command modifies COM parameters that should not need modification under normal operations. Contact EMC Support before using this command.**

**Contexts** All contexts.

**Syntax**

```
connectivity window set
[-n|--name]
{clusterId|minChangeTime|avgWaitUpperBound|
historyDecayDenominator|historyDecayNumerator|
minChangeTime>windowEnlargeDelta>windowSizeInitial
>windowSizeMax>windowSizeMin}
[-v|--value] integer
[-f|--force]
```

**Arguments** **Required arguments**  
[-n|--name] - Name of the value to set. Names include:

**clusterId** - Unique numeric ID assigned to the cluster. Do not attempt to change this value.

**minChangeTime** - The minimum interval in milliseconds between any two com window size calculations. The calculation will cause the history speed to decay, but not necessarily change the window size. High settings slow response to link quality change; low settings cause frequent window size adjustment.  
Default: 1000 msec.

**avgWaitUpperBound** - The threshold, in milliseconds, at which the window size will shrink. If the average waiting time of all outstanding I/Os exceeds this value, the window size will shrink. Otherwise, if full, the window size will grow. Impacts the maximum wait time of the high priority I/Os.  
Default: 1000 msec.

**windowSizeMin** - The minimum com window size.  
Valid range: 1-1400. Default: 1.  
Must be smaller than windowSizeInitial. Setting it too high may cause a longer service wait time.

**windowSizeMax** - The maximum com window size.  
Valid range: 1-1400. Default: 1200.  
Must be larger than windowSizeInitial. Setting it too low may cause low throughput

**windowSizeInitial** - The initial window size when first establishing communication with a peer.  
Valid range: 1-1400. Default: 15.  
Must be larger than windowSizeMin and smaller than windowSizeMax.

**windowEnlargeDelta** - The delta used when increasing the window size. Modify this parameter to change the window size increment speed.

**historyDecayNumerator** - Numerator of the decay factor. See historyDecayDenominator.  
Default: 1.

**historyDecayDenominator** - Denominator of the decay factor:

Decay Factor = historyDecayNumerator/historyDecayDenominator

Higher Decay Factor means history speed has longer and stronger effect on window size; lower means history speed's effect on window size is shorter and weaker.

Default: 2.

**[-v | --value]** - An integer value to assign to the name.

#### Optional arguments

**[-f | --force]** - Force the value to be set, bypassing any confirmations or guards.

- Description** COM bandwidth management senses how fast I/O is sent to a peer and adjusts the I/O window size accordingly. Adjustments ensure that:
- ◆ High priority I/O does not have to wait longer than a expected maximum time before being tried the first time.
  - ◆ Enough I/O is queued to the sending protocol to obtain the best link throughput.

**Example** VPllexcli:/> **connectivity window set windowSizeMin 2**

```
VPllexcli:/> connectivity window show
Name                               Value
-----
avgWaitUpperBound                 1000
historyDecayDenominator            2
historyDecayNumerator              1
minChangeTime                      1000
windowEnlargeDelta                 5
windowSizeInitial                  15
windowSizeMax                      1200
windowSizeMin                       2
```

- See also**
- ◆ [connectivity window show on page 149](#)
  - ◆ [connectivity window stat on page 150](#)

## connectivity window show

Displays values that control the generic operation of communications for one or more directors.

**Contexts** All contexts.

**Syntax** `connectivity window show`

**Description** If the values are the same for all directors, then only one set of values is displayed. Otherwise, the values for each director are displayed individually.

For descriptions of entries in the Name field, see [connectivity window set on page 147](#).

**Example**

```
Vplexcli: /> connectivity window show
Name                               Value
-----
avgWaitUpperBound                  1000
historyDecayDenominator             2
historyDecayNumerator               1
minChangeTime                       1000
windowEnlargeDelta                  5
windowSizeInitial                   15
windowSizeMax                       1200
windowSizeMin                       1
```

**Table 9 connectivity window show field descriptions**

| Field                   | Description                                                                                                                                                                                                                                          |
|-------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| avgWaitUpperBound       | The threshold, in milliseconds, at which the window size will shrink. If the average waiting time of all outstanding I/Os exceeds this value, the window size will shrink. Otherwise, if full, the window size will grow.                            |
| historyDecayDenominator | The history speed information decays over time. The unit speed (used to predict windows size) is: <b>speed = decay-rate * history-speed + (1 - decay-rate) * current-speed</b> where the decay-rate = historyDecayNumerator/historyDecayDenominator. |
| historyDecayNumerator   | Value used to calculate the rate at which speed information decays.                                                                                                                                                                                  |
| minChangeTime           | The minimum interval in milliseconds between any two com window size calculations. The calculation will cause the history speed to decay, but not necessarily change the window size                                                                 |
| windowEnlargeDelta      | The delta of increment each time the window size enlarges.                                                                                                                                                                                           |
| windowSizeInitial       | The initial com window size.                                                                                                                                                                                                                         |
| windowSizeMax           | The maximum com window size.                                                                                                                                                                                                                         |
| windowSizeMin           | The minimum com window size                                                                                                                                                                                                                          |

**See also**

- ◆ [connectivity window set on page 147](#)
- ◆ [connectivity window stat on page 150](#)

## connectivity window stat

Displays the current bandwidth statistics for all directors.

**Syntax**     connectivity window stat

**Contexts**   All contexts.

### Description

### Example

```
VPllexcli:/> connectivity window stat
Statistics for director Cluster_1_Dir1A:
Director          Current Window Size  Outstanding I/O  Queued I/O
-----
Cluster_1_Dir1B   10                   0                0
Cluster_2_Dir_1A  93                   0                0
Cluster_2_Dir_1B  93                   0                0
Cluster_2_Dir_2A  105                  0                0
Cluster_2_Dir_2B  105                  0                0

Statistics for director Cluster_1_Dir1B:
Director          Current Window Size  Outstanding I/O  Queued I/O
-----
Cluster_1_Dir1A   22                   0                0
```

**Table 10**     connectivity window stat field descriptions

| Field                      | Description                                                                                                    |
|----------------------------|----------------------------------------------------------------------------------------------------------------|
| Statistics for director... | Name of the local director.                                                                                    |
| Director                   | Name of the remote director.                                                                                   |
| Current Window Size        | The current window size towards the remote director.                                                           |
| Outstanding I/O            | The number of I/Os currently outstanding (started) to the remote director                                      |
| Queued I/O                 | The number of I/Os in the queue, waiting to start. Outstanding I/O and Queued I/O are exclusive to each other. |

**See also**

- ◆ [connectivity window set on page 147](#)
- ◆ [connectivity window show on page 149](#)

## consistency-group add-virtual-volumes

Adds one or more virtual volumes to a consistency group.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **add-virtual-volumes**

**Syntax**

```
consistency-group add-virtual-volumes
  [-v|--virtual-volumes] virtual-volume, virtual-volume, ...
  [-g|--consistency-group] consistency-group
```

**Arguments** `[-v|--virtual-volumes] virtual-volume, virtual-volume, ...` - \* List of one or more comma-separated glob patterns or context paths of the virtual volume(s) to add.

`[-g|--consistency-group] consistency-group` - \* Context path of the consistency group to which to add the specified virtual volume(s).

If the current context is a consistency-group or is below, then that consistency group is the default. Otherwise, this argument is required.

\* - argument is positional.

**Description** Adds the specified virtual volumes to a consistency group. The properties of the consistency group immediately apply to the added volume.



### IMPORTANT

**Only volumes with visibility and storage-at-cluster properties which match those of the consistency group can be added to the consistency group.**

Maximum # of volumes in a consistency group: 1000

All volumes used by the same application and/or same host should be grouped together in a consistency group.

Only volumes with storage at both clusters (distributed volumes) are allowed in asynchronous consistency groups.

If any of the specified volumes are already in the consistency group, the command skips those volumes, but prints a warning message for each one.



### IMPORTANT

**When adding virtual volumes to a RecoverPoint-enabled consistency group, the RecoverPoint cluster may not note the change for 2 minutes. Wait for 2 minutes between adding virtual volumes to a RecoverPoint-enabled consistency group and creating or changing a RecoverPoint consistency group.**

**Example** Add virtual volumes to a consistency group.

In the following example:

- ◆ The `cd` command changes the context to the target consistency group.
- ◆ The “`consistency-group list-eligible-virtual-volumes`” command displays virtual volumes that are eligible to be added to the consistency group.
- ◆ The `consistency-group add-virtual-volumes` command adds the specified virtual volumes to the consistency group.

- ◆ The `ls` command in displays the change:

```
VPLexcli:/> cd /clusters/cluster-1/consistency-groups/TestCG
```

```
VPLexcli:/clusters/cluster-1/consistency-groups/TestCG> consistency-group
list-eligible-virtual-volumes
[TestDDevice-1_vol, TestDDevice-2_vol, TestDDevice-3_vol, TestDDevice-4_vol,
TestDDevice-5_vol]
```

```
VPLexcli:/clusters/cluster-2/consistency-groups/TestCG> add-virtual-volumes --virtual-volumes
TestDDevice-2_vol
```

**Example** Add multiple volumes using a single command. Enter separate virtual volumes by commas:

```
VPLexcli:/clusters/cluster-1/consistency-groups/TestCG> add-virtual-volumes
TestDDevice-1_vol,TestDDevice-2_vol
```

```
VPLexcli:/clusters/cluster-1/consistency-groups/TestCG> ll
```

Attributes:

| Name                 | Value                                                                                   |
|----------------------|-----------------------------------------------------------------------------------------|
| active-clusters      | []                                                                                      |
| cache-mode           | asynchronous                                                                            |
| detach-rule          | active-cluster-wins                                                                     |
| operational-status   | [(cluster-1,{ summary:: ok, details:: [] }),(cluster-2,{ summary:: ok, details:: [] })] |
| passive-clusters     | [cluster-1, cluster-2]                                                                  |
| recoverpoint-enabled | false                                                                                   |
| storage-at-clusters  | [cluster-1, cluster-2]                                                                  |
| virtual-volumes      | [TestDDevice-1_vol, TestDDevice-2_vol]                                                  |
| visibility           | [cluster-1, cluster-2]                                                                  |

Contexts:

| Name         | Description |
|--------------|-------------|
| advanced     | -           |
| recoverpoint | -           |

- See also**
- ◆ [consistency-group create on page 156](#)
  - ◆ [consistency-group list-eligible-virtual-volumes on page 160](#)
  - ◆ [consistency-group remove-virtual-volumes on page 161](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group choose-winner

Selects a winning cluster during an inter-cluster link failure.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **choose winner**.

**Syntax**

```
consistency-group choose-winner
[-c|--cluster] cluster
[-g|--consistency-group] consistency-group
[-f|--force]
```

**Arguments** `[-c|--cluster] cluster` - \*The cluster on which to roll back and resume I/O.

`[-g|--consistency-group] consistency-group` - \* Context path of the consistency group on which to roll back and resume I/O.

### Optional arguments

`[-f|--force]` - Do not prompt for confirmation. Allows this command to be run using a non-interactive script.

\* - argument is positional.

**Description** Use the **choose-winner** command when:

- ◆ I/O must be resumed on a cluster during a link outage
- ◆ The selected cluster has not yet detached its peer
- ◆ The detach-rules require manual intervention

The selected cluster will detach its peer cluster in preparation for continuing I/O. I/O continues or is suspended depending on the cache mode of the consistency group:

- ◆ For **synchronous** consistency groups: I/O resumes immediately.
- ◆ For **asynchronous** consistency groups: I/O may suspend pending use of the `"consistency-group resume-after-rollback"` command.

For asynchronous consistency groups, it may be required to roll back data to the last consistent image.

Roll-back is required only when both clusters of an asynchronous consistency group were active.



### CAUTION

**When the clusters cannot communicate, it is possible to use this command to select both clusters as the winning cluster (conflicting detach). In a conflicting detach, both clusters resume I/O independently.**

**When the inter-cluster link heals in such a situation, manual intervention is required to pick a winning cluster. The data image of the winning cluster will be used to make the clusters consistent again. Any changes at the losing cluster during the link outage will be discarded.**

**Do not use this command to specify more than one cluster as the winner.**

**Example** Select cluster-2 as the winner for consistency group TestCG:

```
Vplexcli:/clusters/cluster-2/consistency-groups/TestCG> choose-winner --cluster cluster-2
```

WARNING: This can cause data divergence and lead to data loss. Ensure the other cluster is not serving I/O for this consistency group before continuing. Continue? (Yes/No) **Yes**

In the following example:

- ◆ The two **ls** command shows an asynchronous consistency group “cg1” when an inter-cluster link outage has occurred.

The detach-rule is 'no-automatic-winner', so I/O stops at both clusters, the status summary is 'suspended' (showing that I/O has stopped), and the status details contain 'cluster-departure', indicating that I/O has stopped because the clusters can no longer communicate with one another.

- ◆ The **choose winner** command forces cluster-1 to detach cluster-2.
- ◆ The **ls** command displays the change at cluster-1.

Cluster-1 status is 'suspended, requires-resume-after-rollback'.

Cluster-2, is still suspended, cluster-departure.

Cluster-1 is the winner, so it 'detached' cluster-2, but has rolled back its view of the data to the last consistent point\*.

I/O at cluster-1 remains suspended, waiting for the administrator to issue the “**consistency-group resume-after-rollback**” command.

\*Rollback is applicable only to asynchronous consistency groups.

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name                 | Value                                                                                                                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active-clusters      | [cluster-1, cluster-2]                                                                                                                                                                                                                                         |
| cache-mode           | asynchronous                                                                                                                                                                                                                                                   |
| detach-rule          | no-automatic-winner                                                                                                                                                                                                                                            |
| operational-status   | [(cluster-1,{ summary:: suspended, details:: [cluster-departure, rebuilding-across-clusters, restore-link-or-choose-winner] }), (cluster-2,{ summary:: suspended, details:: [cluster-departure, rebuilding-across-clusters, restore-link-or-choose-winner] })] |
| passive-clusters     | []                                                                                                                                                                                                                                                             |
| recoverpoint-enabled | false                                                                                                                                                                                                                                                          |
| storage-at-clusters  | [cluster-1, cluster-2]                                                                                                                                                                                                                                         |
| virtual-volumes      | [dd1_vol, dd2_vol]                                                                                                                                                                                                                                             |
| visibility           | [cluster-1, cluster-2]                                                                                                                                                                                                                                         |

Contexts:

advanced recoverpoint

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> choose-winner -c cluster-1
```

WARNING: This can cause data divergence and lead to data loss. Ensure the other cluster is not serving I/O for this consistency group before continuing. Continue? (Yes/No) **Yes**

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name               | Value                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| active-clusters    | [cluster-1, cluster-2]                                                                                                                                |
| cache-mode         | asynchronous                                                                                                                                          |
| detach-rule        | no-automatic-winner                                                                                                                                   |
| operational-status | [(cluster-1,{ summary:: suspended, details:: [requires-resume-after-rollback] }), (cluster-2,{ summary:: suspended, details:: [cluster-departure] })] |

```
passive-clusters      []  
recoverpoint-enabled false  
storage-at-clusters [cluster-1, cluster-2]  
virtual-volumes      [dd1_vol, dd2_vol]  
visibility            [cluster-1, cluster-2]
```

Contexts:

advanced recoverpoint

- See also**
- ◆ [consistency-group resume-after-data-loss-failure on page 166](#)
  - ◆ [consistency-group resume-after-rollback on page 169](#)
  - ◆ [consistency-group resume-at-loser on page 172](#)
  - ◆ [consistency-group summary on page 178](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group create

Creates and names an empty consistency group.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **create**.

**Syntax**

```
consistency-group create
  [-n|--name] consistency-group name
  [-c|--cluster] cluster
```

**Arguments** `[-n|--name] consistency-group name` - \* Name of the new consistency group. Must be unique on the entire VPLEX.

`[-c|--cluster] cluster` - Context path of the cluster at which to create the consistency group. If the current context is a cluster or below, that cluster is the default. Otherwise, this argument is required.

\* - argument is positional.

**Description** Creates and names an empty consistency group.

A maximum of 1024 consistency groups with synchronous cache mode can be configured.

A maximum of 16 consistency groups with asynchronous cache mode can be configured.

Each consistency group can contain up to 1000 virtual volumes.

All consistency groups have configurable properties that determine I/O behavior, including:

- ◆ cache mode - synchronous or asynchronous. Default is synchronous. Modified using the “set” command.
- ◆ visibility - determines which clusters know about a consistency group. Default is only to the cluster where the consistency group was created. Modified using the “set” command.
- ◆ storage-at-clusters - tells VPLEX at which cluster the physical storage associated with a consistency group is located. Modified using the “set” command.
- ◆ local-read-override - whether the volumes in this consistency group use the local read override optimization. Default is true. Modified using the “set” command.
- ◆ detach-rule - determines the “winning” cluster when there is an inter-cluster link outage. Modified using the “consistency-group set-detach-rule active-cluster-wins”, “consistency-group set-detach-rule no-automatic-winner”, and “consistency-group set-detach-rule winner” commands.
- ◆ auto-resume-at-loser - whether the loser automatically resumes I/O when the inter-cluster link is repaired after a failure. Default is false. Modified using the “set” command in `/clusters/cluster/consistency-groups/consistency-group-name/advanced` context.
- ◆ virtual-volumes - member volumes of the consistency group. Modified using the “consistency-group add-virtual-volumes” and “consistency-group remove-virtual-volumes” commands.
- ◆ recoverpoint-enabled - allows the consistency group to be used in conjunction with RecoverPoint. Default is false. Modified using the “set” command.

**IMPORTANT**

**When enabling or disabling RecoverPoint for a consistency group, the RecoverPoint cluster may not note the change for 2 minutes. Wait for 2 minutes between setting or changing the recoverpoint-enabled property before creating or changing a RecoverPoint consistency group.**

Additional configurable properties are applicable only to consistency groups with asynchronous cache mode, including:

- ◆ default closeout-time - default for the maximum time a delta remains open to accept new writes. Default is 30 seconds. Modified using the “set” command
- ◆ maximum-queue-depth - configures the maximum possible depth of the delta processing queues. Default is 6. Modified using the “set” command.

Refer to the *VPLEX Administration Guide* for more information about the consistency group properties.

**Example** In the following example,

- ◆ The **ls /clusters/\*/consistency-groups/** command displays the names of all consistency groups in both clusters.
- ◆ The **consistency-group create** command creates an empty synchronous consistency group “TestCG”.
- ◆ The **ls** command in consistency group context displays the new name. The **ls TestCG** command displays details about the new consistency group.

**Note:** Refer to the *VPLEX Administration Guide* for a description of the fields in the following examples.

```
Vplexcli:/> ls /clusters/*/consistency-groups/

/clusters/cluster-1/consistency-groups:
test10 test11 test12 test13 test14
test15 test16 test5 test6 test7 test8
test9 vs_RAM_c1wins vs_RAM_c2wins vs_oban005 vs_sun190

/clusters/cluster-2/consistency-groups:
.
.
.
Vplexcli:/> cd /clusters/cluster-1/consistency-groups/

Vplexcli:/clusters/cluster-1/consistency-groups> consistency-group create --name TestCG
--cluster cluster-1

Vplexcli:/clusters/cluster-1/consistency-groups> ls
TestCG test10 test11 test12 test13
test14 test15 test16 test5 test6
test7 test8 test9 vs_RAM_c1wins vs_RAM_c2wins
vs_oban005 vs_sun190

Vplexcli:/clusters/cluster-1/consistency-groups> ls TestCG

/clusters/cluster-1/consistency-groups/TestCG:

Attributes:
Name Value
-----
active-clusters []
cache-mode synchronous
```

```
detach-rule                -
operational-status        [(cluster-1,{ summary:: ok, details:: [] })]
passive-clusters          []
recoverpoint-enabled      true
storage-at-clusters       []
virtual-volumes           []
visibility                 [cluster-1]
```

## Contexts:

```
advanced recoverpoint
```

- See also**
- ◆ [consistency-group add-virtual-volumes on page 151](#)
  - ◆ [consistency-group destroy on page 159](#)
  - ◆ [consistency-group remove-virtual-volumes on page 161](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group destroy

Destroys the specified empty consistency group(s).

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **destroy**.

**Syntax** `consistency-group destroy`  
`[-g|--consistency-group] consistency-group, consistency-group, ...`  
`--force`

**Arguments** **Required arguments**

`[-g|--consistency-group] consistency-group, consistency-group, ...` - \* List of one or more comma-separated context paths of the consistency group(s) to destroy.

**Optional arguments**

`[-f|--force]` - Force the operation to continue without confirmation. Allows this command to be run using a non-interactive script.

\* - argument is positional.

**Description** Destroys the specified consistency group(s).

All clusters where the consistency group is visible must be operational in order for the consistency group to be destroyed.

All clusters where the consistency group has storage-at-clusters must be operational in order for the consistency group to be destroyed.



### WARNING

*Before using the consistency-group destroy command on an asynchronous consistency group, ensure that no host applications are using any member volumes of the consistency group.*

**Example** Destroy the specified consistency group:

```
Vplexcli:/clusters/cluster-1/consistency-groups> destroy TestCG
WARNING: The following items will be destroyed:
Context
-----
/clusters/cluster-1/consistency-groups/TestCG

Do you wish to proceed? (Yes/No)
```

- See also**
- ◆ [consistency-group create on page 156](#)
  - ◆ [consistency-group remove-virtual-volumes on page 161](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group list-eligible-virtual-volumes

Displays the virtual volumes that are eligible to be added to a specified consistency group.

**Contexts** All contexts.

**Syntax** `consistency-group list-eligible-volumes  
[-g|consistency-group] consistency-group`

**Arguments** `[-g|--consistency-group] consistency-group` - The consistency group for which the eligible virtual volumes shall be listed.

If the current context is a consistency group or is below, a consistency group, that consistency group is the default.

Otherwise, this argument is required.

**Description** Displays eligible virtual volumes that can be added to a consistency group. Eligible virtual volumes:

- ◆ Must not be a logging volume
- ◆ Have storage at every cluster in the 'storage-at-clusters' property of the target consistency group
- ◆ Are not members of any other consistency group
- ◆ Have no properties (detach rules, auto-resume) that conflict with those of the consistency group. That is, detach and resume properties of either the virtual volume or the consistency group must not be set.

**Example** List eligible virtual volumes from the target consistency group context:

```
Vplexcli:/clusters/cluster-1/consistency-groups/TestCG2>
list-eligible-virtual-volumes
[dr1_C12_0000_vol, dr1_C12_0001_vol, dr1_C12_0002_vol,
 dr1_C12_0003_vol, dr1_C12_0004_vol, dr1_C12_0005_vol,
 dr1_C12_0006_vol, dr1_C12_0007_vol, dr1_C12_0008_vol,
 dr1_C12_0009_vol, dr1_C12_0010_vol, dr1_C12_0011_vol,
 dr1_C12_0012_vol, dr1_C12_0013_vol, dr1_C12_0014_vol,
 dr1_C12_0015_vol, dgc_p2z_test_vol, vmax_DR1_C1_r1_0000_12_vol,
 vmax_DR1_C1_r0_0000_12_vol,
.
.
.
```

List eligible virtual volumes from the root context:

```
Vplexcli:/> consistency-group list-eligible-virtual-volumes
/clusters/cluster-1/consistency-groups/TestCG2
[dr1_C12_0000_vol, dr1_C12_0001_vol,
 dr1_C12_0002_vol, dr1_C12_0003_vol, dr1_C12_0004_vol,
.
.
.
```

**See also** [consistency-group add-virtual-volumes on page 151](#)  
[consistency-group remove-virtual-volumes on page 161](#)  
[consistency-group summary on page 178](#)

- ◆ *EMC VPLEX Administrator's Guide*

## consistency-group remove-virtual-volumes

Removes one or more virtual volumes from the consistency group.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **remove-virtual-volumes**.

**Syntax**

```
consistency-group remove-virtual-volumes
  [-v|--virtual-volumes] virtual-volume, virtual-volume, ...
  [-g|--consistency-group context path
  --force
```

**Arguments** **Required arguments**

**[-v|--virtual-volumes] *virtual-volume, virtual-volume, ...*** - \*Glob pattern or a list of one or more comma-separated context paths of the virtual volume(s) to remove from the consistency group.

**[-g|--consistency-group] *context path*** - \* Context path of the consistency group from which to remove the specified virtual volume.

If the current context is a consistency-group or is below, then that consistency group is the default. Otherwise, this argument is required.

**--force** - Do not ask for confirmation when removing volumes from an asynchronous consistency group in a VPLEX Geo system. Allows this command to be run using a non-interactive script.

\* - argument is positional.

**Description** Removes one or more virtual volumes from the consistency group.

If the pattern given to `--virtual-volumes` argument matches volumes that are not in the consistency group, the command skips those volumes, and prints a warning message for each one.



### **WARNING**

*Before using the consistency-group remove-virtual-volumes command on an asynchronous consistency group, ensure that no host applications are using the volumes you are removing.*

If this is an asynchronous consistency group in a VPLEX Geo system, and one or more of the volumes being removed is in a storage view, this command asks for confirmation before proceeding. Volumes removed from an asynchronous group become synchronous, and an application that is sensitive to the increased I/O latency may experience data loss or data unavailability.

Best practice is to either:

- ◆ Remove the volumes from the view, or
- ◆ Perform the operation when I/O loads are light.

Use the **--force** argument to suppress the request for confirmation.

**IMPORTANT**

When removing virtual volumes from a RecoverPoint-enabled consistency group, the RecoverPoint cluster may not note the change for 2 minutes. Wait for 2 minutes between removing virtual volumes from a RecoverPoint-enabled consistency group and creating or changing a RecoverPoint consistency group.

**Example** In the following example:

- ◆ The **ls** command displays the virtual volumes in consistency group TestCG.
- ◆ The **consistency-group remove-virtual-volumes** command removes a specified volume from the consistency group.
- ◆ The **ls** command displays the change.

```
VPLEXcli:/> ls /clusters/cluster-1/consistency-groups/TestCG
```

```
/clusters/cluster-1/consistency-groups/TestCG:
```

```
-----
.
.
.
virtual-volumes      [dr1_C12_0919_vol, dr1_C12_0920_vol,
                    dr1_C12_0921_vol, dr1_C12_0922_vol]
visibility            [cluster-1, cluster-2]
.
.
.
```

```
VPLEXcli:/> consistency-group remove-virtual-volumes
/clusters/cluster-1/virtual-volumes/dr1_C12_0920_vol --consistency-group
/clusters/cluster-1/consistency-groups/TestCG
```

```
VPLEXcli:/> ls /clusters/cluster-1/consistency-groups/TestCG
```

```
/clusters/cluster-1/consistency-groups/TestCG:
```

```
-----
Name                    Value
-----
.
.
.
storage-at-clusters     [cluster-1, cluster-2]
synchronous-on-director-failure -
virtual-volumes        [dr1_C12_0919_vol, dr1_C12_0921_vol,
                    dr1_C12_0922_vol]
.
.
.
```

- See also**
- ◆ [consistency-group create on page 156](#)
  - ◆ [consistency-group destroy on page 159](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group resolve-conflicting-detach

Select a winning cluster on a consistency group on which there has been a conflicting detach.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **resolve-conflicting-detach**.

**Syntax**

```
consistency-group resolve-conflicting-detach
  [-c|--cluster] cluster
  [-g|--consistency-group] consistency-group
  [-f|--force]
```

**Arguments** **Required arguments**

`[-c|--cluster] cluster` - \* The cluster whose data image will be used as the source for resynchronizing the data images on both clusters.

`[-g|--consistency-group] consistency-group` - \* The consistency group on which to resolve the conflicting detach.

**Optional arguments**

`[-f|--force]` - Do not prompt for confirmation. Allows this command to be run using a non-interactive script.

\* - argument is positional.

**Description**



**CAUTION**

**This command results in data loss at the losing cluster.**

During an inter-cluster link failure, an administrator may permit I/O to continue at both clusters. When I/O continues at both clusters:

- ◆ The data images at the clusters diverge.
- ◆ Legs of distributed volumes are logically separate.

When the inter-cluster link is restored, the clusters learn that I/O has proceeded independently.

I/O continues at both clusters until the administrator picks a “winning” cluster whose data image will be used as the source to resynchronize the data images.

Use this command to pick the winning cluster. For the distributed volumes in the consistency group:

- ◆ I/O at the “losing” cluster is suspended (there is an impending data change)
- ◆ The administrator stops applications running at the losing cluster.
- ◆ Any dirty cache data at the losing cluster is discarded
- ◆ The legs of distributed volumes rebuild, using the legs at the winning cluster as the rebuild source.

When the applications at the losing cluster are shut down, use the “[consistency-group resume-after-data-loss-failure](#)” command to allow the system to service I/O at that cluster again.

**Example** Select cluster-1 as the winning cluster for consistency group "TestCG" from the TestCG context:

```
Vplexcli:/clusters/cluster-1/consistency-groups/TestCG> resolve-conflicting-detach
This will cause I/O to suspend at clusters in conflict with cluster cluster-1, allowing you to
stop applications at those clusters. Continue? (Yes/No) yes
```

Select cluster-1 as the winning cluster for consistency group "TestCG" from the root context:

```
Vplexcli:/> consistency-group resolve-conflicting-detach --cluster cluster-1
--consistency-group /clusters/cluster-1/consistency-groups/TestCG
This will cause I/O to suspend at clusters in conflict with cluster cluster-1, allowing you to
stop applications at those clusters. Continue? (Yes/No) Yes
```

In the following example, I/O has resumed at both clusters during an inter-cluster link outage. When the inter-cluster link is restored, the two clusters will come back into contact and learn that they have each detached the other and carried on I/O.

- ◆ The **ls** command shows the operational-status as 'ok, requires-resolve-conflicting-detach' at both clusters.
- ◆ The **resolve-conflicting-detach** command selects cluster-1 as the winner. Cluster-2 will have its view of the data discarded. I/O is suspended on cluster-2.
- ◆ The **ls** command displays the change in operational status.
  - At cluster-1, I/O continues, and the status is 'ok'.
  - At cluster-2, the view of data has changed and so I/O is suspended pending the "**consistency-group resume-at-loser**" command.

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name                 | Value                                                                                                                                                          |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active-clusters      | [cluster-1, cluster-2]                                                                                                                                         |
| cache-mode           | asynchronous                                                                                                                                                   |
| detach-rule          | no-automatic-winner                                                                                                                                            |
| operational-status   | [(cluster-1,{ summary:: ok, details:: [requires-resolve-conflicting-detach] }), (cluster-2,{ summary:: ok, details:: [requires-resolve-conflicting-detach] })] |
| passive-clusters     | []                                                                                                                                                             |
| recoverpoint-enabled | false                                                                                                                                                          |
| storage-at-clusters  | [cluster-1, cluster-2]                                                                                                                                         |
| virtual-volumes      | [dd1_vol, dd2_vol]                                                                                                                                             |
| visibility           | [cluster-1, cluster-2]                                                                                                                                         |

Contexts:

advanced recoverpoint

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> resolve-conflicting-detach -c cluster-1
This will cause I/O to suspend at clusters in conflict with cluster cluster-1, allowing you to
stop applications at those clusters. Continue? (Yes/No) Yes
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name            | Value                  |
|-----------------|------------------------|
| active-clusters | [cluster-1, cluster-2] |
| cache-mode      | asynchronous           |
| detach-rule     | no-automatic-winner    |

```
operational-status  [(cluster-1,{ summary:: ok, details:: [] }),  
                    (cluster-2,{ summary:: suspended, details:: [requires-resume-at-loser] })]  
passive-clusters    []  
recoverpoint-enabled false  
storage-at-clusters [cluster-1, cluster-2]  
virtual-volumes     [ddl_vol, dd2_vol]  
visibility           [cluster-1, cluster-2]
```

Contexts:

advanced recoverpoint

- See also**
- ◆ [consistency-group resume-after-data-loss-failure on page 166](#)
  - ◆ [consistency-group resume-at-loser on page 172](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group resume-after-data-loss-failure

Resumes I/O on an asynchronous consistency group when there are data loss failures.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **resume-after-data-loss**.

**Syntax**

```
consistency-group resume-after-data-loss-failure
  [-c|--cluster] cluster
  [-g|--consistency-group] context path
  [-f|--force]
```

**Arguments** **Required arguments**

`[-c|--cluster] cluster` - \* The winning cluster, whose data image is used as the base for resuming I/O.

`[-g|--consistency-group] context path` - \* The consistency group on which to resolve the data loss failure.

**Optional arguments**

`[-f|--force]` - Do not prompt for confirmation. Allows this command to be run using a non-interactive script.

\* - argument is positional.

**Description**

In the event of multiple near-simultaneous director failures, or a director failure followed very quickly by an inter-cluster link failure, an asynchronous consistency group may experience data loss. I/O automatically suspends on the volumes in the consistency group at all participating clusters.

Use this command to resume I/O. Specifically, this command:

- ◆ Selects a winning cluster whose current data image will be used as the base from which to continue I/O.
- ◆ On the losing cluster, synchronizes the data image with the data image on the winning cluster.
- ◆ Resumes I/O at both clusters.

This command may make the data loss larger, because dirty data at the losing cluster may be discarded.

All the clusters participating in the consistency group must be present in order to use this command.

If there has been no data-loss failure in the group, this command prints an error message and does nothing.

**Example**

Select the leg at cluster-1 as the source image to synchronize data for virtual volumes in consistency group "TestCG" (from root context):

```
Vplexcli:/> consistency-group resume-after-data-loss-failure --cluster cluster-1
--consistency-group /clusters/cluster-1/consistency-groups/TestCG
```

Data may be discarded at clusters other than cluster-1. Continue? (Yes/No) **Yes**

Perform the same task from the specific consistency group context:

```
Vplexcli:/clusters/cluster-1/consistency-groups/TestCG> resume-after-data-loss-failure
--cluster cluster-1
```

Data may be discarded at clusters other than cluster-1. Continue? (Yes/No) **Yes**

In the following example, the **auto-resume-at-loser** property of the consistency group is set to true; that is I/O automatically resumes on the losing cluster when connectivity is restored.

- ◆ The **ls** command displays the operational status of a consistency group where cluster-2 is the winner (**detach-rule** is winner) after multiple failures at the same time have caused a data loss. The cluster-1 side of the consistency group is suspended.
- ◆ The **resume-after-data-loss-failure** command selects cluster-1 as the source image from which to re-synchronize data.
- ◆ After a short wait, the **ls** command displays that the cluster-1 side of the consistency group is no longer suspended:

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG1> ls
```

```
Attributes:
Name          Value
-----
active-clusters [cluster-1, cluster-2]
cache-mode     synchronous
detach-rule    winner cluster-2 after 5s
operational-status [suspended, requires-resume-after-data-loss-failure]
passive-clusters []
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes DR1_RAM_c2win_lr1_36_vol, DR1_RAM_c2win_lr1_16_vol,
DR1_RAM_c2win_lr0_6_vol, DR1_RAM_c2win_lrC_46_vol,
.
.
.
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG1> resume-after-data-loss-failure -f -c
cluster-1
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG1> ls
```

```
Attributes:
Name          Value
-----
active-clusters [cluster-2]
cache-mode     synchronous
detach-rule    winner cluster-2 after 5s
operational-status [ok]
passive-clusters [cluster-1]
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes DR1_RAM_c2win_lr1_36_vol, DR1_RAM_c2win_lr1_16_vol,
DR1_RAM_c2win_lr0_6_vol, DR1_RAM_c2win_lrC_46_vol,
.
.
.
```

In the following example, the **auto-resume-at-loser** property of the consistency group is set to false; that is I/O remains suspended on the losing cluster when connectivity is restored. I/O must be manually resumed.

- ◆ The **ls** command displays the operational status of a consistency group where cluster-1 is the winner (**detach-rule** is winner) after multiple failures at the same time have caused a data loss. The cluster-1 side of the consistency group is suspended.

- ◆ The **resume-after-data-loss-failure** command selects cluster-2 as the source image from which to re-synchronize data.
- ◆ After a short wait, the **ls** command displays that the cluster-1 side of the consistency group remains suspended:

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG2> ls
```

```
Attributes:
Name          Value
-----
active-clusters    []
cache-mode        synchronous
detach-rule       winner cluster-1 after 5s
operational-status [suspended, requires-resume-after-data-loss-failure]
passive-clusters  [cluster-1, cluster-2]
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes   DR1_RAM_c2win_lr1_7_vol, DR1_RAM_c2win_lr0_37_vol,
                    DR1_RAM_c2win_lr0_27_vol, DR1_RAM_c2win_lr1_27_vol,
                    DR1_RAM_c2win_lrC_17_vol, DR1_RAM_c2win_lr0_17_vol,
.
.
.
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG2> resume-after-data-loss-failure -f -c cluster-2
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/CG2> ll
```

```
Attributes:
Name          Value
-----
active-clusters    [cluster-2]
cache-mode        synchronous
detach-rule       winner cluster-1 after 5s
operational-status [suspended]
passive-clusters  [cluster-1]
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes   DR1_RAM_c2win_lr1_7_vol, DR1_RAM_c2win_lr0_37_vol,
                    DR1_RAM_c2win_lr0_27_vol, DR1_RAM_c2win_lr1_27_vol,
                    DR1_RAM_c2win_lrC_17_vol, DR1_RAM_c2win_lr0_17_vol,
.
.
.
```

- See also**
- ◆ [consistency-group resume-after-rollback on page 169](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group resume-after-rollback

Resume I/O to the volumes on the winning cluster in a consistency group after:

- ◆ The losing cluster(s) have been detached, and
- ◆ Data has been rolled back to the last point at which all clusters had a consistent view.



### **WARNING**

*In a Geo configuration, on a cluster that successfully vaulted and unvaulted, the user should contact EMC Engineering for assistance before rolling back the data prior to re-establishing communication with the non-vaulting cluster.*

### **Contexts**

All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **resume-after-rollback**.

### **Syntax**

```
consistency-group resume-after-rollback
[-g|--consistency-group context path
[-f|--force]
```

### **Arguments**

#### **Required arguments**

`[-g|--consistency-group] context path` - \* The consistency group on which to resume I/O at the winning cluster.

#### **Optional arguments**

`[-f|--force]` - Do not prompt for confirmation. Allows this command to be run using a non-interactive script.

\* - argument is positional.

### **Description**



### **CAUTION**

**This command addresses data unavailability by ensuring a data loss.**

This command is part of a two-step recovery procedure to allow I/O to continue in spite of an inter-cluster link failure.

1. Use the “`consistency-group choose-winner`” command to select the winning cluster.
2. Use this command to tell the winning cluster to roll back its data image to the last point where the clusters were known to agree, and then proceed with I/O.

The first step in the recovery procedure can be automated by setting a detach-rule-set

The second step is required only if the losing cluster has been “active” -that is, writing to volumes in the consistency group since the last time the data images were identical at the clusters.

If the losing cluster is active, the distributed cache at the losing cluster contains dirty data, and without that data, the winning cluster's data image is inconsistent. Resuming I/O at the winner requires rolling back the winner's data image to the last point where the clusters agreed.

Applications may experience difficulties if the data changes, so the roll-back and resumption of I/O is not automatic.

The delay gives the administrator the chance to halt applications. The administrator then uses this command start the rollback roll-back in preparation for resuming I/O.

The winning cluster rolls back its data image to the last point at which the clusters had the same data images, and then allows I/O to resume at that cluster.

At the losing cluster, I/O remains suspended.

When the inter-cluster link is restored, I/O remains suspended at the losing cluster, unless the 'auto-resume' flag is set to 'true'.

**Example** Resume I/O on the cluster-2 leg of virtual volumes in consistency group TestCG:

```
Vplexcli:/clusters/cluster-2/consistency-groups/TestCG> resume-after-rollback
```

This will change the view of data at cluster cluster-2, so you should ensure applications are stopped at that cluster. Continue? (Yes/No) **Yes**

In the following example:

- ◆ The “**consistency-group choose-winner**” command declares cluster-1 the winning cluster during an inter-cluster link outage.
- ◆ The **ls** command displays a consistency group.

The consistency group is asynchronous, so I/O remains suspended, and requires rollback to the last point where data on the two clusters was the same.

- ◆ The **resume-after-rollback** command rolls back the data image on cluster-1 to the last point where the clusters were known to agree, and then proceeds with I/O.
- ◆ The **ls** command shows that the operational status for the consistency group on cluster-1 is now 'ok'. On cluster-2, the operational status is still suspended, cluster-departure.

When the inter-cluster link heals, cluster-2 will discover that cluster-1 continued servicing I/O without it. Cluster-2 must change its view of data. Operational status is 'suspended, requires-resume-at-loser'

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> choose-winner -c cluster-1
```

WARNING: This can cause data divergence and lead to data loss. Ensure the other cluster is not serving I/O for this consistency group before continuing. Continue? (Yes/No) **Yes**

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name                 | Value                                                                                                                                                                                                         |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active-clusters      | [cluster-1, cluster-2]                                                                                                                                                                                        |
| cache-mode           | asynchronous                                                                                                                                                                                                  |
| detach-rule          | no-automatic-winner                                                                                                                                                                                           |
| operational-status   | [cluster-departure, rebuilding-across-clusters, restore-link-or-choose-winner] }, (cluster-2, { summary:: suspended, details:: [cluster-departure, rebuilding-across-clusters, restore-link-or-choose-winner] |
| passive-clusters     | []                                                                                                                                                                                                            |
| recoverpoint-enabled | false                                                                                                                                                                                                         |
| storage-at-clusters  | [cluster-1, cluster-2]                                                                                                                                                                                        |
| virtual-volumes      | [dd1_vol, dd2_vol]                                                                                                                                                                                            |
| visibility           | [cluster-1, cluster-2]                                                                                                                                                                                        |

Contexts:

advanced recoverpoint

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> resume-after-rollback
```

This will change the view of data at cluster cluster-1, so you should ensure applications are stopped at that cluster. Continue? (Yes/No) **Yes**

```
VPllexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

Attributes:

| Name                 | Value                                                                                                                |
|----------------------|----------------------------------------------------------------------------------------------------------------------|
| active-clusters      | [cluster-1, cluster-2]                                                                                               |
| cache-mode           | asynchronous                                                                                                         |
| detach-rule          | no-automatic-winner                                                                                                  |
| operational-status   | [(cluster-1,{ summary:: ok, details:: [] } ),<br>(cluster-2,{ summary:: suspended, details:: [cluster-departure] })] |
| passive-clusters     | []                                                                                                                   |
| recoverpoint-enabled | false                                                                                                                |
| storage-at-clusters  | [cluster-1, cluster-2]                                                                                               |
| virtual-volumes      | [dd1_vol, dd2_vol]                                                                                                   |
| visibility           | [cluster-1, cluster-2]                                                                                               |

Contexts:

advanced recoverpoint

- See also**
- ◆ [consistency-group choose-winner on page 153](#)
  - ◆ [consistency-group resume-at-loser on page 172](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group resume-at-loser

If I/O is suspended due to a data change, resumes I/O at the specified cluster and consistency group.

**Contexts** All contexts (at the losing cluster).

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **resume-at-loser**.

**Syntax**

```
consistency-group resume-at-loser
[-c|--cluster] cluster
[-s|--consistency-group] consistency-group
[-f|--force]
```

**Arguments** **Required arguments**

`[-c|--cluster] cluster` - \* The cluster on which to roll back and resume I/O.

`[-g|--consistency-group] consistency-group` - \* The consistency group on which to resynchronize and resume I/O.

**Optional arguments**

`[-f|--force]` - Do not prompt for confirmation. Without this argument, the command asks for confirmation to proceed. This protects against accidental use while applications are still running at the losing cluster which could cause applications to misbehave. Allows the command to be executed from a non-interactive script.

\* - argument is positional.

**Description**

During an inter-cluster link failure, an administrator may permit I/O to resume at one of the two clusters: the “winning” cluster.

I/O remains suspended on the “losing” cluster.

When the inter-cluster link heals, the winning and losing clusters re-connect, and the losing cluster discovers that the winning cluster has resumed I/O without it.

Unless explicitly configured otherwise (using the **auto-resume-at-loser** property), I/O remains suspended on the losing cluster. This prevents applications at the losing cluster from experiencing a spontaneous data change.

The delay allows the administrator to shut down applications.

After stopping the applications, the administrator can use this command to:

- ◆ Resynchronize the data image on the losing cluster with the data image on the winning cluster,
- ◆ Resume servicing I/O operations.

The administrator may then safely restart the applications at the losing cluster.

Without the `--force` option, this command asks for confirmation to proceed, since its accidental use while applications are still running at the losing cluster could cause applications to misbehave.

**Example**

```
Vplexcli:/clusters/cluster-2/consistency-groups/TestCG> resume-at-loser
```

This may change the view of data presented to applications at cluster `cluster-2`. You should first stop applications at that cluster. Continue? (Yes/No) **Yes**

In the following example:

- ◆ The **ls** command shows consistency group cg1 as 'suspended, requires-resume-at-loser' on cluster-2 after cluster-2 is declared the losing cluster during an inter-cluster link failure.
- ◆ The **resume-at-loser** command restarts I/O on cluster-2.
- ◆ The **ls** command displays the change in operational status:

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

```
Attributes:
Name          Value
-----
active-clusters [cluster-1, cluster-2]
cache-mode     asynchronous
detach-rule    no-automatic-winner
operational-status [(cluster-1,{ summary:: ok, details:: [] } ),
                  (cluster-2,{ summary:: suspended, details:: [requires-resume-at-loser] })]
passive-clusters []
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes [dd1_vol, dd2_vol]
visibility      [cluster-1, cluster-2]

Contexts:
advanced recoverpoint
```

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> resume-at-loser -c cluster-2
```

This may change the view of data presented to applications at cluster cluster-2. You should first stop applications at that cluster. Continue? (Yes/No) **Yes**

```
Vplexcli:/clusters/cluster-1/consistency-groups/cg1> ls
```

```
Attributes:
Name          Value
-----
active-clusters [cluster-1, cluster-2]
cache-mode     asynchronous
detach-rule    no-automatic-winner
operational-status [(cluster-1,{ summary:: ok, details:: [] } ),
                  (cluster-2,{ summary:: ok, details:: [] } )]
passive-clusters []
recoverpoint-enabled false
storage-at-clusters [cluster-1, cluster-2]
virtual-volumes [dd1_vol, dd2_vol]
visibility      [cluster-1, cluster-2]

Contexts:
advanced recoverpoint
```

- See also**
- ◆ [consistency-group choose-winner on page 153](#)
  - ◆ [consistency-group resume-after-rollback on page 169](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group set-detach-rule active-cluster-wins

Sets the detach-rule for one or more asynchronous consistency groups to “active-cluster-wins”.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **set-detach-rule active-cluster-wins**.

**Syntax** `consistency-group set-detach-rule active-cluster-wins [-g|--consistency-group] consistency-group, consistency-group, ...`

**Arguments** `[-g|--consistency-group] consistency-group, consistency-group, ...` - The consistency group(s) on which to apply the active-cluster-wins detach rule.

**Description** Applies the “active-cluster-wins” detach rule to one or more specified consistency groups.

This detach rule dictates that when only one cluster is doing writes (the 'active' cluster), that cluster wins. Applicable only to asynchronous consistency groups.

Detach rules define the policy for when a cluster may automatically detach its peer cluster in the event of an inter-cluster link failure.

The applicable detach rule varies depending on the visibility, storage-at-clusters, and cache mode properties of the consistency group. [Table 11](#) lists the applicable detach rules for various combinations of visibility, storage-at-clusters, and cache-mode.

**Table 11 Consistency group detach rules**

| Visibility              | Storage-at-clusters     | Cache-mode   | Applicable detach rules                                      |
|-------------------------|-------------------------|--------------|--------------------------------------------------------------|
| cluster-1               | cluster-1               | synchronous  | N/A                                                          |
| cluster-1 and cluster-2 | cluster-1 and cluster-2 | synchronous  | no-automatic-winner,<br>winner cluster-1<br>winner cluster-2 |
| cluster-1 and cluster-2 | cluster-1               | synchronous  | no-automatic-winner,<br>winner cluster-1                     |
| cluster-1 and cluster-2 | cluster-1 and cluster-2 | asynchronous | no-automatic-winner<br>active-cluster-wins                   |



### IMPORTANT

**When RecoverPoint is deployed, it may take up to 2 minutes for the RecoverPoint cluster to take note of changes to a VPLEX consistency group. Wait for 2 minutes after changing the detach rule for a VPLEX consistency group before creating or changing a RecoverPoint consistency group.**

**Example** Set the detach-rule for a single asynchronous consistency group from the group's context:

```
Vplexcli:/clusters/cluster-1/consistency-groups/TestCG> set-detach-rule active-cluster-wins
```

Set the detach-rule for two asynchronous consistency groups from the root context:

```
Vplexcli:/> consistency-group set-detach-rule active-cluster-wins -g /clusters/cluster-1/consistency-groups/TestCG, /clusters/cluster-1/consistency -groups/TestCG2
```

- See also**
- ◆ [consistency-group set-detach-rule no-automatic-winner](#) on page 176
  - ◆ [consistency-group set-detach-rule winner](#) on page 177
  - ◆ *EMC VPLEX Administrator's Guide*

## consistency-group set-detach-rule no-automatic-winner

Sets the detach-rule for one or more asynchronous consistency groups to “no-automatic-winner”.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **set-detach-rule no-automatic-winner**.

**Syntax** `consistency-group set-detach-rule no-automatic-winner [-g|--consistency-group] consistency-group, consistency-group, ...`

**Arguments** `[-g|--consistency-group] consistency-group, consistency-group, ...` - The consistency group(s) on which to apply the no-automatic-winner detach rule.

**Description** Applies the “no-automatic-winner” detach rule to one or more specified consistency groups.

This detach rule dictates no automatic detaches occur in the event of an inter-cluster link failure.

Applicable to both asynchronous and synchronous consistency groups.



### **IMPORTANT**

**When RecoverPoint is deployed, it may take up to 2 minutes for the RecoverPoint cluster to take note of changes to a VPLEX consistency group. Wait for 2 minutes after changing the detach rule for a VPLEX consistency group before creating or changing a RecoverPoint consistency group.**

See Table 11, “Consistency group detach rules,”

**Example** Set the detach-rule for a single consistency group from the group’s context:

```
VPLexcli:/clusters/cluster-1/consistency-groups/TestCG> set-detach-rule no-automatic-winner
```

Set the detach-rule for two consistency groups from the root context:

```
VPLexcli: /> consistency-group set-detach-rule no-automatic-winner -g /clusters/cluster-1/consistency-groups/TestCG, /clusters/cluster-1/consistency-groups/TestCG2
```

**See also**

- ◆ [consistency-group set-detach-rule active-cluster-wins on page 174](#)
- ◆ [consistency-group set-detach-rule winner on page 177](#)
- ◆ *EMC VPLEX Administrator’s Guide*

## consistency-group set-detach-rule winner

Sets the detach-rule for one or more synchronous consistency groups to “winner”.

**Contexts** All contexts.

In `/clusters/cluster-n/consistency-groups/group-name` context, command is **set-detach-rule winner**.

**Syntax**

```
consistency-group set-detach-rule winner
[-c|--cluster] cluster-id
[-d|--delay] seconds
[-g|--consistency-group] consistency-group,
consistency-group, ...
```

**Required arguments** `[-c|--cluster] cluster-id` - The cluster that will be the winner in the event of an inter-cluster link failure.

`[-d|--delay] seconds` - The number of seconds after an inter-cluster link fails before the winning cluster detaches. Valid values for the delay timer are:

**0** - Detach occurs immediately after the link failure is detected.

**number** - Detach occurs after the specified number of seconds have elapsed. There is no practical limit to the number of seconds, but delays longer than 30 seconds won't allow I/O to resume quickly enough to avoid problems with most host applications.

**Optional arguments** `[-g|--consistency-group] consistency-group, consistency-group, ...` - The consistency group(s) on which to apply the winner detach rule.

**Description** Applies the “winner” detach rule to one or more specified synchronous consistency groups.

This detach rule dictates that in the event of an inter-cluster link failure, the cluster specified as the winner detaches after the number of seconds specified by the delay argument.

Applicable only to synchronous consistency groups.



### IMPORTANT

**When RecoverPoint is deployed, it may take up to 2 minutes for the RecoverPoint cluster to take note of changes to a VPLEX consistency group. Wait for 2 minutes after changing the detach rule for a VPLEX consistency group before creating or changing a RecoverPoint consistency group.**

See Table 11, “Consistency group detach rules,”

**Example** Set the detach-rule for a single consistency group from the group's context:

```
VPLexcli:/clusters/cluster-1/consistency-groups/TestCG> set-detach-rule winner --cluster
cluster-1 --delay 5s
```

Set the detach-rule for two consistency groups from the root context:

```
VPLexcli:/> consistency-group set-detach-rule winner --cluster cluster-1 --delay 5s
--consistency-groups TestCG, TestCG2
```

**See also**

- ◆ [consistency-group set-detach-rule active-cluster-wins on page 174](#)
- ◆ [consistency-group set-detach-rule no-automatic-winner on page 176](#)
- ◆ *EMC VPLEX Administrator's Guide*

## consistency-group summary

Displays a summary of all the consistency groups with a state other than 'OK'.

**Contexts** All contexts.

**Syntax** `consistency-group summary`

**Description** Displays all the consistency groups with a state other than 'OK' and the consistency groups at the risk of a rollback.

**Example** Display a summary of unhealthy consistency groups:

```
Vplexcli:/> consistency-group summary
1 consistency groups have status not 'OK'.
```

| Consistency Goup<br>Name | Operational Status | Status<br>Detailed | Active<br>Cluster | Passive<br>Cluster |
|--------------------------|--------------------|--------------------|-------------------|--------------------|
| GRP3                     | Suspended          |                    |                   |                    |

```
0 consistency groups have risk of roll back.
```

**See also**

- [consistency-group create on page 156](#)
- [consistency-group destroy on page 159](#)
- ◆ *EMC VPLEX Administrator's Guide*

---

## date

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

**Contexts** All contexts.

**Syntax** date

**Example** Vplexcli: /> **date**  
Tue Jul 20 15:57:55 UTC 2010

## describe

Describes the attributes of the given context.

**Contexts** All contexts with attributes.

**Syntax** describe  
[*context path*]

**Arguments** **Required arguments**

None.

**Optional arguments**

[-c | --context] *context path* - Context to describe.

**Example** In the following example, the **ll** command displays information about a port, and the **describe** command with no arguments displays additional information.

```
VPLEXcli:/clusters/cluster-2/exports/ports/P000000003CB001CB-B1-FC01> ll
Name                               Value
-----
director-id                        0x000000003cb001cb
discovered-initiators              []
.
```

```
VPLEXcli:/clusters/cluster-2/exports/ports/P000000003CB001CB-B1-FC01> describe
Attribute                           Description
-----
director-id                          The ID of the director where the port is exported.
discovered-initiators                 List of all initiator-ports visible from this port.
.
```

Use the **describe *context*** command to display information about the specified context:

```
VPLEXcli:/> describe --context /clusters/cluster-2/exports/ports/P000000003CB001CB-B1-FC01
Attribute                               Description
-----
/clusters/cluster-2/exports/ports/P000000003CB001CB-B1-FC01::director-id  The ID of the director where the
port is exported.
/clusters/cluster-2/exports/ports/P000000003CB001CB-B1-FC01::discovered-initiators  List of all initiator-ports
visible from this port.
.
```

## device attach-mirror

Attaches a mirror as a RAID 1 child to another (parent) device, and starts a rebuild to synchronize the mirror.

**Contexts** All contexts.

**Syntax**

```
device attach-mirror
  [-d|--device] {context path|name}
  [-m|--mirror] {context path|name}
  [-r|--rule-set] rule-set
  [-f|--force]
```

### Arguments Required arguments

**[-d|--device] context path or name** - \* Name or context path of the device to which to attach the mirror. Does not have to be a top-level device. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

**[-m|--mirror] context path or name** - \* Name or context path of the device to be attached as the mirror. Must be a top-level device. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

### Optional arguments

**[-r|--rule-set] rule-set** - Rule-set to apply to the distributed device that is created when a mirror is added to a local device.

**[-f|--force]** - Do not ask for confirmation when attaching a mirror. Allows this command to be run using a non-interactive script.

If the **--force** argument is not used, prompts for confirmation in two circumstances when the mirror is remote and the parent device must be transformed into a distributed device:

- ◆ The rule set that will be applied to the new distributed device potentially allows conflicting detaches; and
- ◆ In VPLEX GEO systems, where the parent device already has a volume in a storage view. The new distributed device will be synchronous. Applications using the volume may experience greater I/O latency than on the original local device. If an application is sensitive to this latency, it may experience data unavailability.

\* - argument is positional.

**Description** If the parent device is a RAID 0 or RAID c, it is converted to RAID 1.

If the parent device and mirror device are from different clusters, a distributed-device is created.

A storage-volume extent cannot be used as a mirror if the parent device is a distributed-device, or if the parent device is at a different cluster than the storage-volume extent.

If the **--rule-set** argument is omitted, VPLEX assigns a default rule-set to the distributed device as follows:

- ◆ If the parent device has a volume, the distributed device inherits the rule-set of the (exported) parent.
- ◆ If the parent device does not have a volume, the cluster that is local to the management server is assumed to be the winner.

Once determined, VPLEX displays a notice as to which rule-set the created distributed-device has been assigned.

**Note:** If the RAID 1 device is added to a consistency group, the consistency group's detach rule overrides the device's detach rule.

Use the “[rebuild status](#)” command to display the rebuild's progress.

**Example** Attach a mirror without specifying a rule-set (allow VPLEX to select the rule-set):

```
Vplexcli:/clusters/cluster-1/devices> virtual-volume create test_r0c_1  
Vplexcli:/clusters/cluster-1/devices> device attach-mirror --device test_r0c_1 --mirror  
test_r0c_2  
Distributed device 'regression_r0c_1' is using rule-set 'cluster-1-detaches'.
```

Attach a mirror:

```
Vplexcli:/> device attach-mirror --device /clusters/cluster-1/devices/site1device0 --mirror  
/clusters/cluster-1/devices/site1mirror
```

- See also**
- ◆ [consistency-group set-detach-rule winner on page 177](#)
  - ◆ [device detach-mirror on page 184](#)
  - ◆ [local-device create on page 279](#)
  - ◆ [rebuild status on page 366](#)

---

## device collapse

Collapses a one-legged device until a device with two or more children is reached.

**Contexts** All contexts.

**Syntax** `device collapse [-d|--device] context [path|name]`

**Arguments** **Required arguments**  
`[-d|--device] context path or name` - \* Name or context path of the device to collapse.

Does not have to be a top-level device. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

\* - argument is positional.

**Description** If a RAID 1 device is left with only a single child (after removing other children), use the **device collapse** command to collapse the remaining structure. For example:

If RAID 1 device "A" has two child RAID 1 devices "B" and "C", and child device "C" is removed, A is now a one-legged device, but with an extra layer of abstraction:

```
  A
  |
  B
  .. / \ ..
```

Use **device collapse** to remove this extra layer, and change the structure into:

```
  A
  .. / \ ..
```

Applicable to one-legged devices that are not top-level.

**Example** `Vplexcli:/clusters/cluster-1/devices> device collapse --device /clusters/cluster-1/devices/A`

**See also** ♦ [local-device create on page 279](#)

## device detach-mirror

Detaches an up-to-date child from the parent RAID 1 or removes the child from a parent RAID 1.

**Contexts** All contexts.

**Syntax**

```
device detach-mirror
  [-d|--device] [context path|name]
  [-m|--mirror] [context path|name]
  [-r|--rule-set] context path
  [-s|--slot] slot number
  [-i|--discard]
  [-f|--force]
```

**Arguments** **Required arguments**

**[-d|--device] context path or name** - \* Name or context path of the device from which to detach the mirror. Does not have to be a top-level device. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

**[-m|--mirror] context path or name** - \* Name or context path of the mirror to detach. Does not have to be a top-level device. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

**Optional arguments**

**[-s|--slot] slot number** - Slot number of the mirror to be discarded. Applicable only when the **--discard** argument is used.

**[-i|--discard]** - When specified, discards the data in the mirror that is to be detached.

**[-f|--force]** - Force the mirror to be discarded. Must be used when the **--discard** argument is used.

\* - argument is positional.

**Description** If the **--discard** argument is not used, this command detaches an up-to-date child from the parent RAID 1. The detached child becomes a top level device with the same virtual volumes as its parent.

If the **--discard** argument is used, this command removes the child from a parent RAID 1. The detached child will become a top level device but without any virtual volumes or guarantees about its data.

**Example**

```
Vplexcli:/clusters/cluster-1> device detach-mirror --device
/clusters/cluster-1/devices/cluster1device0 --mirror
/clusters/cluster-1/devices/cluster8mirror
```

Detach a “dead” mirror leg from a distributed device.

In the following example:

- ◆ The **ll** command in `/distributed-storage/distributed-devices` context displays a stressed distributed device.
- ◆ The **ll device-name/distributed-device-components** command displays the components of the device (output is truncated)  
Note the Slot number of the failed leg in the display.
- ◆ The **device detach-mirror** command removes the failed device using the slot number displayed in the previous step.

- ◆ The `ll` command in `/distributed-storage/distributed-devices` context displays the change:

```
VPLEXcli:/distributed-storage/distributed-devices> ll
Name                               Status  Operational  Health State  Auto  Rule  WOF  Transfer
-----  -----  -----  -----  -----  ---  ---  ---  -----
-----  -----  -----  -----  -----  ---  ---  ---  -----
ESX_stretched_device               running ok           ok           true   colin -     2M
bbv_temp_device                     running ok           ok           true   colin -     2M
dd_source_device                   running ok           ok           true   colin -     2M
ddt                                 running ok           ok           true   colin -     2M
dev_test_dead_leg_2                 running stressed    major-failure -     colin -     2M
windows_big_drive                   running ok           ok           true   colin -     2M
.
.
.
```

```
VPLEXcli:/distributed-storage/distributed-devices> ll
/dev_test_dead_leg_2_DD/distributed-device-components/

/distributed-storage/distributed-devices/dev_test_dead_leg_2_DD/distributed-device-components
:
```

```
Name                               Cluster  Child  Fully  Operational  Health ..
-----  -----  ---  ---  -----  -----  ---
-----  -----  ---  ---  -----  -----  ---
dev_test_alive_leg_1                cluster-1 1    true   ok           ok
dev_test_dead_leg_2                 cluster-2 0    true   error        critical-failure
```

```
VPLEXcli:/distributed-storage/distributed-devices> device detach-mirror --slot 0 --discard
--force --device /distributed-storage/distributed-devices/dev_test_dead_leg_2
```

```
VPLEXcli:/distributed-storage/distributed-devices> ll
Name                               Status  Operational  Health  Auto  Rule  WOF  Transfer
-----  -----  -----  -----  -----  ---  ---  ---  -----
-----  -----  -----  -----  -----  ---  ---  ---  -----
ESX_stretched_device               running ok           ok           true   colin -     2M
bbv_temp_device                     running ok           ok           true   colin -     2M
dd_source_device                   running ok           ok           true   colin -     2M
ddt                                 running ok           ok           true   colin -     2M
dev_test_dead_leg_2_DD              running ok           ok           -     colin -     2M
windows_big_drive                   running ok           ok           true   colin -     2M
.
.
.
```

**See also** ◆ [device attach-mirror on page 181](#)

## device resume-link-down

Resumes I/O for devices on the winning island during a link outage.

**Contexts** All contexts.

**Syntax**

```
device resume-link-down
  [-c|--cluster] context path
  [-r|--devices] [context path|name]
  [-a|--all-at-island]
  [-f|--force]
```

**Arguments** **Required arguments**

### Optional arguments

**[-c|--cluster] context path** - Resume I/O on the specified cluster and the clusters it is in communication with during a link outage.

Applicable only when the **all-at-island** argument is used or when the specified devices are distributed devices.

Not required for local devices with global visibility.

**[-r|--devices] context path or name** - Name or context path of the devices for which to resume I/O. They must be top-level devices.

**[-a|--all-at-island]** - Resume I/O on all devices on the chosen winning cluster and the clusters with which it is communicating.

**[-f|--force]** - Force the I/O to resume.

### Description

Used when the inter-cluster link fails. Allows one or more suspended mirror legs to resume I/O immediately.

For example, used when the peer cluster is the winning cluster but is known to have failed completely.

Resumes I/O on the specified cluster and the clusters it is in communication with during a link outage.

Detaches distributed devices from those clusters that are not in contact with the specified cluster or detaches local devices from those clusters that are not in contact with the local cluster.



### WARNING

*The device resume-link-down command causes I/O to resume on the local cluster regardless of any rule-sets applied to the device. Verify that rules and any manual detaches do not result in conflicting detaches (cluster-1 detaching cluster-2, and cluster-2 detaching cluster-1). Conflicting detaches will result in lost data on the losing cluster, a full rebuild, and degraded access during the time of the full rebuild.*

When the inter-cluster link fails in a VPLEX Metro or Geo configuration, distributed devices are suspended at one or more clusters. When the rule-set timer expires, the affected cluster is detached.

Alternatively, use the **device resume-link-down** command to detach the cluster immediately without waiting for the rule-set timer to expire.



## WARNING

*Verify that rules and any manual detaches do not result in conflicting detaches (cluster-1 detaching cluster-2, and cluster-2 detaching cluster-1). Conflicting detaches will result in lost data on the losing cluster, a full rebuild, and degraded access during the time of the full rebuild.*

Only one cluster should be allowed to continue for each distributed device. Different distributed devices can have different clusters continue.

Use the `ll /distributed-storage/distributed-devices/device` command to display the rule set applied to the specified device.

Use the `ll /distributed-storage/rule-sets/rule-set/rules` command to display the detach timer for the specified rule-set.

### Example

```
Vplexcli:/distributed-storage/distributed-devices> device  
  resume-link-down --all-at-island --cluster --devices DD_5d --force
```

### See also

- ◆ [device resume-link-up on page 188](#)
- ◆ [ds dd declare-winner on page 222](#)

## device resume-link-up

Resumes I/O on suspended top level devices, virtual volumes, or all virtual volumes in the VPLEX.

**Contexts** All contexts.

**Syntax**

```
device resume-link-up
  [-r|--devices] context path,context path...
  [-v|--virtual-volumes] context path,context path...
  [-a|--all]
  [-f|--force]
```

**Arguments** [-r|--devices] *context path, context path...* - List of one or more context paths or names of the devices for which to resume I/O. They must be top-level devices. If the device name is used, verify that the name is unique throughout the VPLEX, including local devices on other clusters.

[-v|--virtual-volume] *context path,context path...* - Resume I/O on the specified virtual volumes.

[-a|--all] -Resume I/O on all virtual volumes on the losing cluster.

[-f|--force] - Force the I/O to resume.

**Description** Use this command after a failed link is restored, but I/O is suspended at one or more clusters.

Usually applied to the mirror leg on the losing cluster when auto-resume is set to false.

During a WAN link outage, after cluster detach, the primary cluster detaches to resume operation on the distributed device.

If the auto-resume property of a remote or distributed device is set to false and the link has come back up, use the **device resume-link-up** command to manually resume the second cluster.

**Example** Resume I/O on two specified devices:

```
VPlexcli:/distributed-storage/distributed-devices> device
resume-link-up --devices CLAR0014_LUN17_1, CLAR0014_LUN18_1 --force
```

Resume I/O on a specified virtual volume:

```
VPlexcli:/> device resume-link-up --virtual-volumes
/clusters/cluster-1/virtual-volumes/ESX_DataStore1_vol --force
```

Resume I/O on all virtual volumes in the losing cluster:

```
VPlexcli:/> device resume-link-up --all --force
```

**See also** ♦ [device resume-link-down on page 186](#)

## director appcon

Runs the application console on Linux systems.

**Contexts** All contexts.

**Syntax**

```
director appcon
  --xterm-opts options
  [-t|--targets] target glob,target glob...
  --timeout seconds
  --show-plan-only
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**--xterm-opts** -Additional options for xterm on Linux.

Default: -sb -sl 2000.

**[-t|--targets] *target-glob, target-glob...*** - List of one or more glob patterns. Operates on the specified targets. Globs may be a full path glob, or a name pattern. If only a name pattern is supplied, the command finds allowed targets whose names match. Entries must be separated by commas.

Omit this argument if the current context is at or below the target context.

**--timeout *seconds*** - Sets the command timeout. Timeout occurs after the specified number of seconds multiplied by the number of targets found.

Default: 180 seconds per target.

0: No timeout.

**--show-plan-only** - Shows the targets that will be affected, but the actual operation is not performed. Recommended when the **--targets** argument is used.

**Description** Applicable only to Linux systems.

Opens the hardware application console for each director in a separate window.

**Example** Display the available targets:

```
Vplexcli: /> director appcon --show-plan-only

For /engines/engine-2-1/directors/Cluster_2_Dir_1B:

Planned target.
For /engines/engine-2-1/directors/Cluster_2_Dir_1A:

Planned target.
For /engines/engine-2-2/directors/Cluster_2_Dir_2A:
.
.
.
```

**See also** ♦ [director appstatus on page 192](#)

## director appdump

Downloads an application dump from one or more boards.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **appdump**.

**Syntax**

```
director appdump
  [-d|--dir] directory
  [-s|--no-timestamp]
  [-o|--overwrite]
  [-c|--include-cores]
  [-z|--no-zip]
  [-p|--no-progress]
  [-t|--targets] target glob,target glob...
  --timeout seconds
  --show-plan-only
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**[-d|--dir] *directory*** - Directory in which to put dump files.

Default: Current directory.

**[-s|--no-timestamp]** - Omits the timestamp in the dump filenames.

**[-o|--overwrite]** - Overwrites existing files.

Default: Not to overwrite.

**[-c|--include-cores]** - Includes the director's core image in the dump.



### **WARNING**

*Including core images can make this command take a very long time.*

**[-z|--no-zip]** - Turns off the packaging of dump files into a compressed zip file.

**[-p|--no-progress]** - Turns off progress reporting.

**[-t|--targets] *target-glob, target-glob...*** - List of one or more glob patterns. Operates on the specified targets. Globs may be a full path glob or a name pattern. If only a name pattern is supplied, the command finds allowed targets whose names match. Entries must be separated by commas.

Omit this argument if the current context is at or below the target context.

**--timeout *seconds*** - Sets the command timeout. Timeout occurs after the specified number of seconds multiplied by the number of targets found.

Default: 180 seconds per target.

0: No timeout.

**--show-plan-only** - Shows the targets that will be affected, but the actual operation is not performed. Recommended when the **--targets** argument is used.

**Description** Used by automated scripts and by EMC Customer Support to help troubleshoot problems.

The hardware name and a timestamp are embedded in the dump filename. By default, the name of the dump file is:

---

hardware name-YYYY.MM.DD-hh.mm.ss.zip.

**Note:** Timeout is automatically set to 0 (infinite) when dumping core.

**Example** Show the targets available for the appdump procedure:

```
Vplexcli:/> director appdump --show-plan-only
```

```
For /engines/engine-2-1/directors/dirB:
```

```
Planned target.
```

```
For /engines/engine-2-1/directors/dirA:
```

```
Planned target.
```

```
For /engines/engine-1-1/directors/DirA:
```

```
Planned target.
```

```
For /engines/engine-1-1/directors/DirB:
```

```
Planned target.
```

- See Also**
- ◆ [cluster configdump on page 81](#)
  - ◆ [collect-diagnostics on page 106](#)
  - ◆ [getsysinfo on page 272](#)
  - ◆ [sms dump on page 448](#)

## director appstatus

Displays the status of the application on one or more boards.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **appstatus**.

**Syntax**

```
director appstatus
[-t|--targets] target glob, target glob...
--timeout seconds
--show-plan-only
```

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-t|--targets] target-glob, target-glob...` - List of one or more glob patterns. Operates on the specified targets. Globs may be a full path glob or a name pattern. If only a name pattern is supplied, the command finds allowed targets whose names match. Entries must be separated by commas.

Omit this argument if the current context is at or below the target context.

`--timeout seconds` - Sets the command timeout. Timeout occurs after the specified number of seconds multiplied by the number of targets found.

Default: 180 seconds per target.

0: No timeout.

`--show-plan-only` - Shows the targets that will be affected, but the actual operation is not performed. Recommended when the `--targets` argument is used.

**Description** Used by automated scripts and by EMC Customer Support to help troubleshoot problems.

**Example** `VPlexcli:/engines/engine-1-1/directors> appstatus`

For `/engines/engine-1-1/directors/Cluster_1_Dir1B:`

| Application          | Status  | Details |
|----------------------|---------|---------|
| 00601610672e201522-2 | running | -       |

For `/engines/engine-1-1/directors/Cluster_1_Dir1A:`

| Application         | Status  | Details |
|---------------------|---------|---------|
| 00601610428f20415-2 | running | -       |

**See also**

- ◆ [director appcon on page 189](#)
- ◆ [director appdump on page 190](#)

---

## director commission

Starts the director's participation in the cluster.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **commission**.

**Syntax**

```
director commission
[-n|--director] director
[-f|--force]
[-t|--timeout seconds]
[-a|--apply-cluster-settings]
```

**Arguments** **Required arguments**

`[-n|--director] director` - \* The director to be commissioned.

**Optional arguments**

`[-f|--force]` - Commission the director regardless of firmware version mismatch.

`--timeout seconds` - The maximum time to wait for `--apply-cluster-settings` operations to complete, in seconds.

Default: 60 seconds.

0: No timeout.

`[a|--apply-cluster-settings]` - Add this director to a running cluster and apply any cluster-specific settings. Use this argument when adding or replacing a director in an existing VPLEX.

\* - argument is positional.

**Description**

In order to participate in a cluster, a director must be explicitly commissioned. Uncommissioned directors can boot but do not participate in any cluster activities.

Use the **version -a** command to display the firmware version for all directors in the cluster.

The **director commission** command fails if the director's firmware version is different than the already commissioned directors, unless the `--force` argument is used.

**Example** Add a director to a running cluster using the default timeout (60 seconds):

```
Vplexcli:/> director commission --director Cluster_1_Dir1A --apply-cluster-settings
```

**See also**

- ◆ [director decommission on page 194](#)
- ◆ [version on page 501](#)

---

## director decommission

Decommissions a director. The director stops participating in cluster activities.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **decommission**.

**Syntax** `director decommission  
[-n|--director] director`

**Arguments** **Required arguments**  
`[-n|--director] director` - The director to de-commission.

**Description** This command removes the director from participating in the VPLEX, and initializes it to only a partial operational state. The director is no longer a replication target and its front-end ports are disabled.

Then it reboots the director.

**Example** `VPlexcli:/> director decommission --director Cluster_1_Dir1A`

**See also**

- ◆ [director commission on page 193](#)
- ◆ [director forget on page 198](#)
- ◆ [director shutdown on page 201](#)

## director fc-port-stats

Displays/resets Fibre Channel port statistics for a specific director.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **fc-port-stats** *director*.

In `/engines/engine/directors/director` context, command is **fc-port-stats**

**Syntax**

```
director fc-port-stats
[-d|--director] director
[-o|--role] role
[-r|--reset]
```

**Arguments** **Required arguments**

`[-d|--director] director` - Context path of the director for which to display FC statistics. Not required if the current context is `/engines/engine/directors/director`.

**Optional arguments**

`[-o|--role] role` - Filter the ports included in the reply by their role. If no role is specified, all ports at the director are included. This argument is ignored if `--reset` is specified. Roles include:

- **back-end** - The port is used to access storage devices that the system itself does I/O to.
- **front-end** - The port is used to make storage available to hosts.
- **inter-director-communication** - The port is used to communicate with other directors.
- **local-com** - The port is used to communicate with other directors at the same cluster.
- **management** - The port is used to communicate with the management server.
- **wan-com** - The port is used to communicate with other clusters.

`[-r|--reset]` - Reset the statistics counters of all ports at the specified director. If this argument is specified, `--role` is ignored.

**Description** Displays statistics generated by the Tachyon driver for FibreChannel ports at the specified director and optionally with the specified role, or reset said statistics' Run this command from the `/engines/engine/directors/director` context to display the Fibre Channel statistics for the director in the current context.

**Example** Display a director's Fibre Channel port statistics from the root context:

```
Vplexcli: /> director fc-port-stats -d director-2-1-A
```

**Example** Reset the port statistics counters on a director's Fibre Channel ports from the root context:

```
Vplexcli: /> director fc-port-stats -d director-2-1-A --reset
```

**Example** Display a director's Fibre Channel port statistics from the *director* context:

```
Vplexcli: /engines/engine-1-1/directors/director-1-1-A> fc-port-stat
```

Results for director 'director-2-1-A' at Fri Feb 10 16:10:15 MST 2012:

|              |         |         |         |         |         |         |
|--------------|---------|---------|---------|---------|---------|---------|
| Port:        | A1-FC00 | A1-FC01 | A1-FC02 | A1-FC03 | A3-FC00 | A3-FC01 |
| Frames:      |         |         |         |         |         |         |
| - Discarded: | 0       | 0       | 0       | 0       | 0       | 0       |

```

- Expired:          0          0          0          0          0          0
- Bad CRCs:         0          0          0          0          0          0
- Encoding Errors: 0          0          0          0          0          0
Requests:
- Accepted:         0          0          0          0          7437       7437
- Rejected:         0          0          0          0          0          0
- Started:          0          0          0          0          7437       7437
- Completed:        0          0          0          0          7437       7437
- Timed-out:        0          0          0          0          0          0
Tasks:
- Received:         0          0          0          0          7437       7437
- Accepted:         0          0          0          0          7437       7437
- Rejected:         0          0          0          0          0          0
- Started:          0          0          0          0          7437       7437
- Completed:        0          0          0          0          7437       7437
- Dropped:          0          0          0          0          0          0

```

**Table 12** director fc-port-stats field descriptions

| Field                                                                                                                   | Description                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Frames:</b> frames discarded, expired, or with CRC or encoding errors on this port since the counter was last reset. |                                                                                                                                                                             |
| Discarded                                                                                                               | Discarded frames.                                                                                                                                                           |
| Expired                                                                                                                 | Expired frames.                                                                                                                                                             |
| Bad CRCs                                                                                                                | Frames with invalid CRC.                                                                                                                                                    |
| Encoding Errors                                                                                                         | Frames with encoding errors.                                                                                                                                                |
| <b>Requests:</b> outgoing scsi commands issued by a VPLEX initiator.                                                    |                                                                                                                                                                             |
| Accepted                                                                                                                | scsi requests accepted. Requests that have not been rejected.                                                                                                               |
| Rejected                                                                                                                | scsi requests rejected. Requests can be rejected due to: Not an initiator, Port not ready, no memory for allocation.                                                        |
| Started                                                                                                                 | scsi requests started.                                                                                                                                                      |
| Completed                                                                                                               | Requests that have been completed.                                                                                                                                          |
| Timed-out                                                                                                               | Number of requests timed out. Requests have been sent but not responded to within 10 seconds.                                                                               |
| <b>Tasks:</b> incoming scsi commands issued to a VPLEX target.                                                          |                                                                                                                                                                             |
| Received                                                                                                                | Tasks received by the tach driver from the TSDK.                                                                                                                            |
| Accepted                                                                                                                | Tasks started back to the TSDK.                                                                                                                                             |
| Rejected                                                                                                                | Tasks rejected due to: 1. connection not ready 2. IO was aborted by TSDK 3. Tried to send more data than the initiator requested. 4. Unable to allocate enough SGL entries. |
| Started                                                                                                                 | Tasks sent to the upper layer for processing.                                                                                                                               |
| Completed                                                                                                               | Tasks completed to the upper layer.                                                                                                                                         |
| Dropped                                                                                                                 | Tasks dropped due to: 1. Connection not ready. 2. No available task buffers to handle the incoming task.                                                                    |

**See also** ♦ [monitor stat-list on page 325](#)

## director firmware show-banks

Display the status of the two firmware banks for all or specified director(s).

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **firmware show-banks**.

**Syntax** `director firmware show-banks`  
`[-t|--targets] director, director...`

**Arguments** **Required arguments**  
None.

**Optional arguments**

`[-t|--targets] director, director...` - List of one or more names of directors. Display information only for the specified directors. Entries must be separated by commas.

**Description** Show firmware status and version for one or more directors.

**Example** Show firmware banks for two specified directors:

```
Vplexcli:/engines> director firmware show-banks --targets Cluster_1_Dir1A, Cluster_1_Dir1B  
[Director Cluster_1_Dir1B]:
```

```
Banks      Status      Marked for Next Reboot      Director Software Version  
-----  
Bank A     inactive    no                           1.2.43.0.0  
Bank B     active      yes                           1.2.43.2.0
```

```
[Director Cluster_1_Dir1A]:
```

```
Banks      Status      Marked for Next Reboot      Director Software Version  
-----  
Bank A     inactive    no                           1.2.43.0.0  
Bank B     active      yes                           1.2.43.2.0
```

**Table 13** director firmware show-banks field descriptions

| Field                     | Description                                                                                                                                                                           |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Banks                     | Each director has two firmware banks; A and B.                                                                                                                                        |
| Status                    | <b>active</b> - The software in this bank is currently operating on the director.<br><b>inactive</b> - The software in this bank is not operating on the director.                    |
| Marked for next reboot    | <b>no</b> - The software in this bank will not be used the next time the director reboots.<br><b>yes</b> - The software in this bank will be used the next time the director reboots. |
| Director Software version | Software version currently operating in the director.                                                                                                                                 |

**See also** ♦ [version on page 501](#)

## director forget

Removes a director from the VPLEX.

**Contexts** All contexts.

**Syntax** `director forget`  
`[-n|--director] director uuid`

**Arguments** **Required arguments**  
`[-n|--director] director uuid` - Director ID number. Use the **ll** command in `engines/engine/directors` context to display director ID numbers.

**Description** Removes the specified director from the context tree. Deletes all information associated with the director.

**Example** In the following example:

- ◆ The **ll** command in `engines/engine/directors` context displays director IDs.
- ◆ The **director forget** command instructs VPLEX to delete all records pertaining to the specified director.

```
VPlexcli:/engines/engine-1-1/directors> ll
Name                Director ID          Cluster  Commissioned  Operational  Communication
-----            -
Cluster_1_Dir1A    0x000000003ca00147  1        true          ok           ok
Cluster_1_Dir1B    0x000000003cb00147  1        true          ok           ok
```

```
VPlexcli:/engines/engine-1-1/directors> director forget --director 0x000000003ca00147
```

- See also**
- ◆ [director commission on page 193](#)
  - ◆ [director decommission on page 194](#)

---

## director passwd

Changes the access password for the specified director.

**Contexts** All contexts.

In `/engines/engine/directors/director` context, command is **passwd**.

**Syntax**

```
director passwd
  [-n|--director] director
  [-c|--current-password] current-password
  [-p|--new-password] new-password
```

**Arguments** **Required arguments**

**[-n|--director] *director*** - The remote director on which to change the access password.

**[-c|--current-password] *current-password*** - The current access password of the specified director.

**[-p|--new-password] *new-password*** - The new access password to set for the specified director.

**Description** Changes the password for a specified director.

**Example**

**See also** ♦ [user passwd on page 489](#)

## director ping

Displays the round-trip latency from a given director to the target machine, excluding any VPLEX overhead.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **ping**.

**Syntax**

```
director ping
  [-i|--ip-address] ip-address
  [-n|--director] director
  [-w|--wait] [1 - 2147483647]
```

**Arguments** **Required arguments**

`[-i|--ip-address] IP address` - The target's IP address.

**Optional arguments**

`[-n|--director] director` - The director from which to perform the operation.

`[-w|--wait] seconds` - Number of seconds to wait for a response.

Range: 1 - 2147483647

Default: 5.

**Description** ICMP traffic must be permitted between clusters for this command to work properly.

**Notes** To verify that ICMP is enabled, log in to the shell on the management server and use the **ping IP address** command where the IP address is for a director in the VPLEX.

If ICMP is enabled on the specified director, a series of lines is displayed:

```
service@ManagementServer:~> ping 128.221.252.36
PING 128.221.252.36 (128.221.252.36) 56(84) bytes of data.
64 bytes from 128.221.252.36: icmp_seq=1 ttl=63 time=0.638 ms
64 bytes from 128.221.252.36: icmp_seq=2 ttl=63 time=0.591 ms
64 bytes from 128.221.252.36: icmp_seq=3 ttl=63 time=0.495 ms
64 bytes from 128.221.252.36: icmp_seq=4 ttl=63 time=0.401 ms
64 bytes from 128.221.252.36: icmp_seq=5 ttl=63 time=0.552 ms

--- 128.221.252.36 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4002ms
rtt min/avg/max/mdev = 0.401/0.535/0.638/0.084 ms
```

If ICMP is disabled, nothing is displayed.

Press Ctrl-C to exit from ping.

**Example** Ping from root context:

```
VPlexcli:/>director ping -n director-1-1-A -i 192.168.30.67
Round-trip time to 192.168.30.67: 0.111 ms
```

Ping from director context:

```
VPlexcli:/engines/engine-1-1/directors/director-1-1-A>ping 192.168.30.68
Round-trip time to 192.168.30.68: 0.117 ms
```

Remote address is unreachable:

```
VPlexcli:/engines/engine-1-1/directors/director-1-1-A>ping 128.221.252.36
128.221.252.36 is unreachable.
```

**See also** ♦ [director tracepath on page 202](#)

## director shutdown

Starts the orderly shutdown of a director's firmware

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **shutdown**.

**Syntax** `director shutdown`  
`[-f|--force]`  
`[-n|--director] context path`

**Arguments** **Required arguments**  
`[-f|--force]` - Forces this operation.

**Optional arguments**  
`[-n|--director] context path` - \* Director to shut down.  
\* - argument is positional.

**Description** Shuts down the director firmware.

**Note:** Does not shut down the operating system on the director.

After shutdown, state of the director is as follows:

- ◆ Power is on.
- ◆ Director OS running.
- ◆ Director firmware (GeoSynchrony) is stopped.

**Example** In the following example:

- ◆ The **director shutdown** command shuts down DirA.
- ◆ The **ll** command displays the shutdown director.

```
Vplexcli:/engines/engine-1-1/directors/DirA> director shutdown --force  
Please check the status of the director: DirA for its shutdown status.
```

```
Status      Description  
-----  
Started.    Shutdown started.
```

```
Vplexcli:/engines/engine-1-1/directors/DirA> ll
```

```
Attributes:  
Name                                     Value  
-----  
auto-boot                                true  
auto-restart                             true  
.  
.  
marker-led                               off  
operational-status                       stopped  
.  
.
```

- See also**
- ◆ [cluster shutdown on page 86](#)
  - ◆ [director commission on page 193](#)

## director tracepath

Displays the route taken by packets from a specified director to the target machine.

**Contexts** All contexts.

In `/engines/engine/directors` context, command is **tracepath**.

**Syntax**

```
director tracepath
[-i|--ip-address] ip-address
[-n|--director] director
```

**Arguments** **Required arguments**

`[-i|--ip-address] IP address` - The target's IP address. This address is one of the Ethernet WAN ports on another director. Use the **ll port-group** command to display the Ethernet WAN ports on all directors.

**Optional arguments**

`[-n|--director]` - The name of the director from which to perform the operation. Can be either the director's name (for example "director-1-1-A") or an IP address.

**Description** Displays the hops, latency, and MTU along the route from the specified director to the target at the specified IP address.

The number of hops does not always correlate to the number of switches along the route. For example, a switch with a firewall on each side is counted as two hops.

The reported latency at each hop is the round-trip latency from the source hop.

The MTU reported at each hop is limited by the MTU of previous hops and therefore not necessarily the configured MTU at that hop.

In the following illustration, a director in cluster-2 (director-2-1-A) is pinged from a director in cluster-1 (director-1-1-A) in single engine clusters. The ping travels through both management servers before reaching its target:

```
Vplexcli:/> director tracepath -n director-1-1-A/ -i 128.221.252.67
Destination 128.221.252.67 reachable. 3 hops:
Source: endpoint 128.221.253.35, latency 0.049ms, mtu 1500, reachable source
Hop 1: endpoint 128.221.253.33, latency 0.187ms, mtu 1500, reachable mgmt server
Hop 2: endpoint 128.221.253.33, latency 0.176ms, mtu 1428, reachable mgmt server
Hop 3: endpoint 128.221.252.67, latency 1.34ms, mtu 1428, reachable target
```



### CAUTION

**If the target machine does not respond properly, the traceroute may stall. Run this command multiple times.**

Use the **ll /clusters/\*/cluster-connectivity/port-groups** command to display the Ethernet WAN ports on all directors (for both clusters).

**Example** In the following example:

- ◆ The **ll engines/\*/directors** command displays the names of all directors in a 2-cluster VPLEX.
- ◆ The **ll /clusters/\*/cluster-connectivity/port-groups** command displays the IP addresses of the Ethernet WAN ports on all directors.
- ◆ The **director tracepath** command displays the route from the specified director to the specified address using the director's name:

```
Vplexcli:/> ll /clusters/*/cluster-connectivity/port-groups
```

```
/clusters/cluster-1/cluster-connectivity/port-groups:
Name          Subnet          Option Set      Enabled         Member Ports
-----
port-group-0  cluster-1-SN00  optionset-com-0 all-enabled     engine-1-1|director-1-1-A|A2-XG00|192.168.11.140|cluster-1-SN00|optionset-com-0|enabled,
port-group-1  cluster-1-SN01  optionset-com-1 all-enabled     engine-1-1|director-1-1-B|B2-XG00|192.168.11.142|cluster-1-SN00|optionset-com-0|enabled,
engine-1-1|director-1-1-A|A2-XG01|192.168.12.140|cluster-1-SN01|optionset-com-1|enabled,
engine-1-1|director-1-1-B|B2-XG01|192.168.12.142|cluster-1-SN01|optionset-com-1|enabled
```

```
/clusters/cluster-2/cluster-connectivity/port-groups:
Name          Subnet          Option Set      Enabled         Member Ports
-----
port-group-0  cluster-2-SN00  optionset-com-0 all-enabled     engine-2-1|director-2-1-A|A2-XG00|192.168.21.144|cluster-2-SN00|optionset-com-0|enabled,
port-group-1  cluster-2-SN01  optionset-com-1 all-enabled     engine-2-1|director-2-1-B|B2-XG00|192.168.21.146|cluster-2-SN00|optionset-com-0|enabled,
engine-2-1|director-2-1-A|A2-XG01|192.168.22.144|cluster-2-SN01|optionset-com-1|enabled,
engine-2-1|director-2-1-B|B2-XG01|192.168.22.146|cluster-2-SN01|optionset-com-1|enabled
```

```
Vplexcli:/> ll /engines/*/directors
```

```
/engines/engine-1-1/directors:
Name          Director ID      Cluster  Commissioned  Operational  Communication
-----
director-1-1-A 0x0000000044603198 1        true          ok           ok
director-1-1-B 0x0000000044703198 1        true          ok           ok
```

```
/engines/engine-2-1/directors:
Name          Director ID      Cluster  Commissioned  Operational  Communication
-----
director-2-1-A 0x00000000446031b1 2        true          ok           ok
director-2-1-B 0x00000000447031b1 2        true          ok           ok
```

```
Vplexcli:/> director tracepath --director director-1-1-A --ip-address 192.168.22.146
Destination 192.168.22.146 reachable. 2 hops:
Source: endpoint 192.168.12.140, latency 0.029ms, mtu 1500, reachable
Hop 1: endpoint 192.168.12.1, latency 52.6ms, mtu 1500, reachable
Hop 2: endpoint 192.168.22.146, latency 108ms, mtu 1500, reachable
```

Display the route from the specified director to the specified address using the director's IP address:

```
Vplexcli:/> director tracepath 10.6.211.91 -i 10.12.136.12
```

**See also** ♦ [director ping on page 200](#)

---

## director uptime

Prints the uptime information for all connected directors.

**Contexts** All contexts.

In engines/*engine*/directors context, command is **uptime**.

**Syntax** `director uptime`

**Description** Uptime measures the time a machine has been up without any downtime.

**Example**

```
VPLEXcli:/> director uptime  
Director Cluster_1_Dir1B: 42 days, 0 hours, 26 minutes, 17 seconds.  
Director Cluster_1_Dir1A: 42 days, 0 hours, 37 minutes, 46 seconds.  
Director Cluster_2_Dir_1B: 42 days, 0 hours, 25 minutes, 11 seconds.  
Director Cluster_2_Dir_1A: 42 days, 0 hours, 37 minutes, 4 seconds.  
Director Cluster_2_Dir_2B: 42 days, 0 hours, 25 minutes, 37 seconds.  
Director Cluster_2_Dir_2A: 42 days, 0 hours, 37 minutes, 18 seconds.
```

**See also**

- ◆ [cluster summary on page 93](#)
- ◆ [director firmware show-banks on page 197](#)

---

## dirs

Displays the current context stack.

**Contexts** All contexts.

**Syntax** `dirs`

**Description** The stack is displayed from top to bottom, in left to right order.

**Example**

```
Vplexcli:/> dirs
[/]

Vplexcli:/> cd /engines/engine-1-1/

Vplexcli:/engines/engine-1-1> dirs
[/engines/engine-1-1]

Vplexcli:/engines/engine-1-1> cd /directors/

Vplexcli:/engines/engine-1-1/directors> dirs
[/engines/engine-1-1/directors]
```

**See also** ♦ [tree on page 484](#)

## disconnect

Disconnects one or more connected directors.

**Contexts** All contexts.

**Syntax** `disconnect [-n|--directors] context path, context path...`

**Arguments** **Required arguments**  
`[-n|--directors] context path, context path...` - List of one or more remote directors from which to disconnect. Entries must be separated by commas.

**Description** Stops communication from the client to the remote directors and frees up all resources that are associated with the connections.



### **CAUTION**

**Removes the entry in the connections file for the specified director(s).**

**Example** `Vplexcli:> disconnect --directors engines/engine-1-1/directors/DirB`

**See also** ♦ [connect on page 137](#)

## dm migration cancel

Cancels an existing data migration.

**Contexts** All contexts.

In all data-migration (device or extent) contexts, command is **migration cancel**.

In data-migrations/extent-migrations context, command is **cancel**.

**Syntax**

```
dm migration cancel
[-f|--force]
[-m|--migrations] context path,context path...
```

**Arguments** **Required arguments**

**[-m|--migrations] context path,context path...** \* List of one or more migrations to cancel. Entries must be separated by commas.

**Optional arguments**

**[-f|--force]** - Forces the cancellation of the specified migration(s).

\* - argument is positional.

**Description** Use the **dm migration cancel --force --migrations context-path** command to cancel a migration.

Specify the migration by name if that name is unique in the global namespace. Otherwise, specify a full context path.

Migrations can be canceled in the following circumstances:

- ◆ The migration is in progress or paused. The migration is stopped, and any resources it was using are freed.
- ◆ The migration has not been committed. The source and target devices or extents are returned to their pre-migration state.

A migration cannot be canceled if it has been committed.

To remove the migration record from the context tree, see [dm migration remove on page 211](#).

**Example** Cancel a migration from device-migration context:

```
Vplexcli:/data-migrations/device-migrations> dm migration cancel
--force --migrations migrate_012
```

Cancel a device migration from root context:

```
Vplexcli:> dm migration cancel --force --migrations
/data-migrations/device-migrations/migrate_012
```

**See also**

- ◆ [dm migration commit on page 209](#)
- ◆ [dm migration pause on page 210](#)
- ◆ [dm migration remove on page 211](#)
- ◆ [dm migration resume on page 212](#)
- ◆ [dm migration start on page 213](#)

## dm migration clean

Cleans a committed data migration.

**Contexts** All contexts.

In /data-migrations context, command is **migration clean**.

In /data-migrations/device-migrations context, command is **clean**.

In /data-migrations/extent-migrations context, command is **clean**.

**Syntax**

```
dm migration clean
[-f|--force]
[-e|--rename-target]
[-m|--migrations] context path,context path...
```

**Arguments** **Required arguments**

**[-m|--migrations] context path,context path...** \* List of one or more migrations to clean. Entries must be separated by commas.

**Optional arguments**

**[-f|--force]** - Forces the cancellation of the specified migration(s).

**[-e|--rename-target]** - For device migrations only, renames the target device after the source device. If the target device is renamed, the virtual volume on top of it is also renamed if the virtual volume has a system-assigned default name.

\* - argument is positional.

**Description** For device migrations, cleaning dismantles the source device(s) down to its storage volumes. The storage volumes no longer in use are unclaimed.

For device migrations only, use the **--rename-target** argument to rename the target device after the source device. If the target device is renamed, the virtual volume on top of it is also renamed if the virtual volume has a system-assigned default name.

Without renaming, the target devices retain their target names, which can make the relationship between volume and device less evident.

For extent migrations, cleaning destroys the source extent and unclaims the underlying storage-volume if there are no extents on it.

### Example

```
VPlexcli:/data-migrations/device-migrations> dm migration clean --force --migrations
migrate_012
```

Cleaned 1 data migration(s) out of 1 requested migration(s).

- See also**
- ◆ [dm migration cancel on page 207](#)
  - ◆ [dm migration commit on page 209](#)
  - ◆ [dm migration pause on page 210](#)
  - ◆ [dm migration remove on page 211](#)
  - ◆ [dm migration resume on page 212](#)
  - ◆ [dm migration start on page 213](#)

## dm migration commit

Commits a completed data migration allowing for its later removal.

**Contexts** All contexts.  
In /data-migrations context, command is **migration commit**.  
In /data-migrations/extent-migrations context, command is **commit**.  
In /data-migrations/device-migrations context, command is **commit**.

**Syntax** `dm migration commit`  
`[-f|--force]`  
`[-m|--migrations] context path,context path...`

**Arguments** **Required arguments**  
`[-m|--migrations] context path,context path...` \* List of one or more migrations to commit. Entries must be separated by commas.  
`[-f|--force]` - Forces the commitment of the specified migration(s).  
\* - argument is positional.

**Description** The migration process inserts a temporary RAID 1 structure above the source device/extent with the target device/extent as an out-of-date leg of the RAID 1. The migration can be understood as the synchronization of the out-of-date leg (the target).  
After the migration is complete, the commit step detaches the source leg of the RAID 1 and removes the RAID 1.  
The virtual volume, device or extent is identical to the one before the migration except that the source device/extent is replaced with the target device/extent.  
A migration must be committed in order to be cleaned.



### CAUTION

**Verify that the migration has completed successfully before committing the migration.**

**Example** Commit a device migration:  

```
Vplexcli:/data-migrations/device-migrations> commit -m  
beta_device_mig1 -f  
Committed 1 data migration(s) out of 1 requested migration(s).
```

**See also**

- ◆ [dm migration cancel on page 207](#)
- ◆ [dm migration pause on page 210](#)
- ◆ [dm migration remove on page 211](#)
- ◆ [dm migration resume on page 212](#)
- ◆ [dm migration start on page 213](#)

## dm migration pause

Pauses the specified in-progress or queued data migrations.

**Contexts** All contexts.

In /data-migrations context, command is **migration pause**.

In /data-migrations/extent-migrations context, command is **pause**.

In /data-migrations/device-migrations context, command is **pause**.

**Syntax** `dm migration pause  
[-m|--migrations] context path,context path...`

**Arguments** **Required arguments**

`[-m|--migrations] context path,context path...` \* List of one or more migrations to pause. Entries must be separated by commas.

\* - argument is positional.

**Description** Active migrations (a migration that has been started) can be paused and then resumed at a later time.

Pause an active migration to release bandwidth for host I/O during periods of peak traffic.

Specify the migration by name if that name is unique in the global namespace. Otherwise, specify a full pathname.

Use the **dm migration resume** command to resume a paused migration.

**Example** Pause a device migration:

```
Vplexcli:/data-migrations/device-migrations> dm migration pause --migrations migrate_012
```

- See also**
- ◆ [dm migration cancel on page 207](#)
  - ◆ [dm migration commit on page 209](#)
  - ◆ [dm migration remove on page 211](#)
  - ◆ [dm migration resume on page 212](#)
  - ◆ [dm migration start on page 213](#)

---

## dm migration remove

Removes the record of canceled or committed data migrations.

**Contexts** All contexts.

In /data-migrations context, command is **migration remove**.

In /data-migrations/extent-migrations context, command is **remove**.

In /data-migrations/device-migrations context, command is **remove**.

**Syntax** `dm migration remove  
[-f|--force]  
[-m|--migrations] context path,context path...`

**Arguments** **Required arguments**

`[-m|--migrations] context path,context path...` \* List of one or more migrations to remove. Entries must be separated by commas.

`[-f|--force]` - Forces the removal of the specified migration(s).

\* - argument is positional.

**Description** Before a migration record can be removed, it must be canceled or committed to release the resources allocated to the migration.

**Example** Remove a migration:

```
Vplexcli:/data-migrations/device-migrations> remove -m beta_device_mig1 -f
```

```
Removed 1 data migration(s) out of 1 requested migration(s).
```

- See also**
- ◆ [dm migration cancel on page 207](#)
  - ◆ [dm migration commit on page 209](#)
  - ◆ [dm migration pause on page 210](#)
  - ◆ [dm migration resume on page 212](#)
  - ◆ [dm migration start on page 213](#)

## dm migration resume

Resumes a previously paused data migration.

**Contexts** All contexts.

In /data-migrations context, command is **migration resume**.

In/data-migrations/extent-migrations context, command is **resume**.

In/data-migrations/device-migrations context, command is **resume**.

**Syntax** `dm migration resume  
[-m|--migrations] context path,context path...`

**Arguments** **Required arguments**

`[-m|--migrations] context path,context path...`- List of one or more migrations to resume. Entries must be separated by commas.

**Description** Active migrations (a migration that has been started) can be paused and then resumed at a later time.

Pause an active migration to release bandwidth for host I/O during periods of peak traffic.

Use the **dm migration resume** command to resume a paused migration.

**Example** Resume a paused device migration:

```
VPLEXcli:/data-migrations/device-migrations> dm migration resume --migrations migrate_012
```

**See also**

- ◆ [dm migration cancel on page 207](#)
- ◆ [dm migration commit on page 209](#)
- ◆ [dm migration pause on page 210](#)
- ◆ [dm migration remove on page 211](#)
- ◆ [dm migration start on page 213](#)

## dm migration start

Starts the specified migration.

**Contexts** All contexts.

In /data-migrations context, command is **migration start**.

in/data-migrations/extent-migrations context, command is **start**.

in/data-migrations/device-migrations context, command is **start**.

**Syntax**

```
dm migration start
[-n|--name] migration name
[-s|--transfer-size] [40 K - 128 M]
[-f|--from] [source extent|source device]
[-t|--to] [target extent|target device]
--paused
--force
```

**Arguments** **Required arguments**

**[-n|--name] migration name...** - \* Name of the new migration. Used to track the migration's progress, and to manage (cancel, commit, pause, resume) the migration.

**[-f|--from] {source extent | source device}** - \* The name of source extent or device for the migration. Specify the source device or extent by name if that name is unique in the global namespace. Otherwise, specify a full pathname.

If the source is an extent, the target must also be an extent. If the source is a device, the target must also be a device.

**[-t|--to] {target extent | target device}** - \* The name of target extent or device for the migration. Specify the target device or extent by name if that name is unique in the global namespace. Otherwise, specify a full pathname.

**Optional arguments**

**[-s|--transfer-size] value** - Maximum number of bytes to transfer per operation per device. A bigger transfer size means smaller space available for host I/O. Must be a multiple of 4 K.

Range: 40 KB - 128 M. Default: 128 K.

If the host I/O activity is very high, setting a large transfer size may impact host I/O. See [About transfer-size on page 53](#).

**--force** - Do not ask for confirmation. Allows this command to be run using a non-interactive script.

**--paused** - Starts the migration in a paused state.

\* - argument is positional.

**Description** Starts the specified migration. If the target is larger than the source, the extra space on the target is unusable after the migration. If the target is larger than the source, a prompt to confirm the migration is displayed.

Up to 25 local and 25 distributed migrations (rebuids) can be in progress at the same time. Any migrations beyond those limits are queued until an existing migration completes.



### CAUTION

**Migration of volumes in asynchronous consistency groups is not supported on volumes that are in use. Schedule this activity as a maintenance activity to avoid Data Unavailability.**

**Extent migrations** - Extents are ranges of 4K byte blocks on a single LUN presented from a single back-end array. Extent migrations move data between extents in the same cluster. Use extent migration to:

- ◆ Move extents from a “hot” storage volume shared by other busy extents,
- ◆ Defragment a storage volume to create more contiguous free space,
- ◆ Support technology refreshes.

Start and manage extent migrations from the extent migration context:

```
Vplexcli: /> cd /data-migrations/extent-migrations/
```

```
Vplexcli: /data-migrations/extent-migrations>
```

**Device migrations** - Devices are RAID 0, RAID 1, or RAID C built on extents or other devices. Devices can be nested; a distributed RAID 1 can be configured on top of two local RAID 0 devices. Device migrations move data between devices on the same cluster or between devices on different clusters. Use device migration to:

- ◆ Migrate data between dissimilar arrays,
- ◆ Relocate a “hot” volume to a faster array,



### WARNING

*Device migrations are not recommended between clusters. All device migrations are synchronous. If there is I/O to the devices being migrated, and latency to the target cluster is greater than 5ms, significant performance degradation may occur.*

Start and manage device migrations from the device migration context:

```
Vplexcli: /> cd /data-migrations/device-migrations/
```

```
Vplexcli: /data-migrations/device-migrations>
```

### Prerequisites for target devices/extents

The target device or extent of a migration must:

- ◆ Be the same size or larger than the source device or extent  
If the target is larger in size than the source, the extra space cannot be utilized. For example, if the source is 200 GB, and the target is 500 GB, only 200 GB of the target can be used after a migration. The remaining 300 GB cannot be claimed.
- ◆ Not have any existing volumes on it.

**Example** Start a device migration:

```
Vplexcli: /data-migrations/device-migrations> dm migration start --name migrate_012 --from device_012 --to device_012a
```

## Migration procedure overview: device migration

Use the following general steps to perform a device migration.

1. Navigate to the device migration context:

```
Vplexcli: /> cd /data-migrations/device-migrations/
```

```
Vplexcli: /data-migrations/device-migrations>
```

2. Use the **dm migration start** command to start a one-time device migration:

- ◆ Name of migration: `migrate_012`
- ◆ Source device: `device_012`
- ◆ Target device: `device_012a`

```
Vplexcli: /data-migrations/device-migrations> dm migration start --name migrate_012 --from device_012 --to device_012a
```

3. Navigate to the migration's context, and use the **ls** command to monitor the migration's progress:

```
Vplexcli: /data-migrations/device-migrations> cd /migrate_012
```

```
Vplexcli: /data-migrations/device-migrations/migrate_012> ls
```

| Name            | Value                        |
|-----------------|------------------------------|
| from-cluster    | cluster-1                    |
| percentage-done | 10                           |
| source          | device_012                   |
| source-exported | false                        |
| start-time      | Fri May 28 13:32:23 MDT 2010 |
| status          | in progress                  |
| target          | device_012a                  |
| target-exported | false                        |
| to-cluster      | cluster-2                    |
| transfer-size   | 12M                          |

4. Optionally, pause and resume an in-progress or queued migration.

Use the **dm migration pause --migrations migration-name** command to pause a migration.

```
Vplexcli: /data-migrations/device-migrations> dm migration pause --migrations migrate_012
```

Use the **dm migration resume --migrations migration-name** command to resume a paused migration.

```
Vplexcli: /data-migrations/device-migrations> dm migration resume --migrations migrate_012
```

5. When the migration's status changes to "complete", use the **dm migration commit** command to commit the migration.

The migration cannot be canceled after this step.

```
Vplexcli: /data-migrations/device-migrations> dm migration commit --force --migrations migrate_012
```

```
Committed 1 data migration(s) out of 1 requested migration(s).
```

6. Clean the migration.

For device migrations, cleaning dismantles the source device down to its storage volumes. The storage volumes no longer in use are unclaimed.

For extent migrations, cleaning destroys the source extent and unclaims the underlying storage-volume if there are no extents on it.

```
Vplexcli: /data-migrations/device-migrations> dm migration clean --force --migrations migrate_012
```

```
Cleaned 1 data migration(s) out of 1 requested migration(s).
```

7. If necessary, remove the migration record:

```
Vplexcli:/data-migrations/device-migrations> dm migration remove --force --migrations  
migrate_012
```

Removed 1 data migration(s) out of 1 requested migration(s).

- See also**
- ◆ [batch-migrate create-plan on page 48](#)
  - ◆ [batch-migrate start on page 53](#)
  - ◆ [dm migration cancel on page 207](#)
  - ◆ [dm migration commit on page 209](#)
  - ◆ [dm migration pause on page 210](#)
  - ◆ [dm migration remove on page 211](#)
  - ◆ [dm migration resume on page 212](#)

## drill-down

Displays the components of a view, virtual volume or device, down to the storage-volume context.

**Contexts** All contexts.

**Syntax**

```
drill-down
  [-v|--storage-view] context path,context path...
  [-o|--virtual-volume] context path,context path...
  [-r|--device] context path,context path...
```

**Arguments** **Required arguments**

**[-v|--storage-view] context path,context path...** - List of one or more views to drill down. Entries must be separated by commas. Glob style pattern matching is supported.

**[-o|--virtual-volume] context path,context path...** - List of one or more virtual-volumes to drill down. Entries must be separated by commas. Glob style pattern matching is supported.

**[-r|--device] context path,context path...** - List of one or more devices to drill down. Entries must be separated by commas. Glob style pattern matching is supported.

**Description** Displays the components of the specified object.

To display a list of available objects, use the **drill-down *object-type*** command followed by the <TAB> key, where *object type* is storage-view, device, or virtual-volume.

**Example** Display the components of a virtual volume:

```
Vplexcli:/clusters/cluster-2> drill-down --virtual-volume dd_21_vol/
virtual-volume: dd_21_vol (cluster-2)
  distributed-device: dd_21
    distributed-device-component: dev1723_614 (cluster-2)
      extent: extent_Symm1723_614_1
        storage-volume: Symm1723_614 (blocks: 0 - 4195199)
    distributed-device-component: dev1852_214 (cluster-1)
      extent: extent_Symm1852_214_1
        storage-volume: Symm1852_214 (blocks: 0 - 4195199)
```

Display the elements of a storage view:

```
Vplexcli:/clusters/cluster-2> drill-down --storage-view exports/storage-views/LicoJ010
storage-view: LicoJ010 (cluster-2)
  virtual-volume: base01_vol (cluster-2)
    local-device: base01 (cluster-2)
      extent: extent_base_volume_1
        storage-volume: base_volume (blocks: 0 - 524287)
  virtual-volume: dd_00_vol (cluster-2)
    distributed-device: dd_00
      distributed-device-component: dev1723_00 (cluster-1)
        extent: extent_Symm1723_200_1
          storage-volume: Symm1723_200 (blocks: 0 - 4195199)
        extent: extent_Symm1723_204_1
          storage-volume: Symm1723_204 (blocks: 0 - 4195199)
        extent: extent_Symm1723_208_1
          storage-volume: Symm1723_208 (blocks: 0 - 4195199)
        extent: extent_Symm1723_20C_1
          storage-volume: Symm1723_20C (blocks: 0 - 4195199)
```

.

```
.
virtual-volume: dev_Symm1723_91C_vol (cluster-2)
  local-device: dev_Symm1723_91C (cluster-2)
    extent: extent_Symm1723_91C_1
      storage-volume: Symm1723_91C (blocks: 0 - 4195199)
  iport: LicoJ010_hba1
  iport: LicoJ010_hba0
  iport: LicoJ010_hba3
  iport: LicoJ010_hba2
  tport: P000000003CB000E6-B1-FC00
  tport: P000000003CB001CB-B0-FC00
  tport: P000000003CA001CB-A0-FC00
  tport: P000000003CA000E6-A0-FC00
  tport: P000000003CB001CB-B1-FC00
  tport: P000000003CB000E6-B0-FC00
  tport: P000000003CA001CB-A1-FC00
  tport: P000000003CA000E6-A1-FC0
```

Display the components of a device:

```
VPlexcli:/clusters/cluster-2/devices> drill-down --device dev_Symm1723_918
local-device: dev_Symm1723_918 (cluster-2)
  extent: extent_Symm1723_918_1
    storage-volume: Symm1723_918 (blocks: 0 - 4195199)
```

**See also** ♦ [tree on page 484](#)

## ds dd create

Creates a new distributed-device.

**Contexts** All contexts.

**Syntax**

```
ds dd create
  [-n|name] name
  [-d|--devices context path,context path,...]
  [-l|--logging-volumes context path,context path,...]
  [-r|rule-set] rule-set
  [-f|--force]
```

**Arguments** **Required arguments**

**[-n|--name] name** - \* The name of the new distributed-device. Must be unique across the VPLEX.

**[-d|--devices] context path, context path...** - \* List of one or more local devices that will be legs in the new distributed-device.

**[-l|--logging-volume] context path, context path...** - List of one or more logging volumes to use with this device. If no logging volume is specified, a logging volume is automatically selected from any available logging volume that has sufficient space for the required entries. If no available logging volume exists, an error message is returned.

\* - argument is positional.

**Optional arguments**

**[-r|--rule-set]** - The rule-set to apply to the new distributed device. If the **--rule-set** argument is omitted, the cluster that is local to the management server is assumed to be the winner in the event of an inter-cluster link failure.

**[-f|--force]** - Forces a rule-set with a potential conflict to be applied to the new distributed device.

**Description** The new distributed device consists two “legs”; local devices on each cluster.



**WARNING**

*A device created by this command does not initialize its legs, or synchronize the contents of the legs. Because of this, consecutive reads of the same block may return different results for blocks that have never been written. Host reads at different clusters are almost certain to return different results for the same unwritten block, unless the legs already contain the same data.*



**CAUTION**

**Use this command only if the resulting device will be initialized using tools on the host.**

**Do not use this command if one leg of the resulting device contains data that must be preserved. Applications using the device may corrupt the pre-existing data.**

To create a device when one leg of the device contains data that must be preserved, use the “[device attach-mirror](#)” command to add a mirror to the leg. The data on the leg will be copied automatically to the new mirror.

The individual local devices may include any underlying type of storage volume or geometry (RAID 0, RAID 1, or RAID C), but they should be the same capacity.

If a distributed device is configured with local devices of different capacities:

- ◆ The resulting distributed device will be only as large as the smaller local device.
- ◆ The “leftover” capacity on the larger device will *not* be available.

To create a distributed device without wasting capacity, choose local devices on each cluster with the same capacity.

The geometry of the new device is automatically 'RAID 1'.

Each cluster in the VPLEX may contribute a maximum of one component device to the new distributed device.



### CAUTION

If there is pre-existing data on a storage-volume, and the storage-volume is not claimed as being *application consistent*, converting an existing local RAID device to a distributed RAID using the `ds dd create` command will *not* initiate a rebuild to copy the data to the other leg. Data will exist at only one cluster. To prevent this, do one of the following:

- 1) Claim the disk with data using the application consistent flag, or
- 2) Create a single-legged RAID 1 or RAID 0 and add a leg using the device `attach-mirror` command.

Use the “`set`” command to enable/disable automatic rebuilds on the distributed device. The rebuild setting is immediately applied to the device.

- ◆ `set rebuild-allowed true` starts or resumes a rebuild if mirror legs are out of sync.
- ◆ `set rebuild-allowed false` stops a rebuild in progress.

When set to true, the rebuild continues from the point where the it was halted. Only those portions of the device that have not been rebuilt are affected. The rebuild does not start over.

**Example** In the following example, the `ds dd create` command creates a new distributed device with the following attributes:

- ◆ Name: ExchangeDD
- ◆ Devices:
  - `/clusters/cluster-2/devices/s6_exchange`
  - `/clusters/cluster-1/devices/s8_exchange`
- ◆ Logging volumes:
  - `/clusters/cluster-1/system-volumes/cluster_1_loggingvol`
  - `/clusters/cluster-2/system-volumes/cluster_2_loggingvol`
- ◆ Rule-set: rule-set-7a

```
VPLexcli:/distributed-storage/distributed-devices> ds dd create --name ExchangeDD --devices
/clusters/cluster-2/devices/s6_exchange,/clusters/cluster-1/devices/s8_exchange
--logging-volumes
/clusters/cluster-1/system-volumes/cluster_1_loggingvol,/clusters/cluster-2/system-volumes/cl
uster_2_loggingvol --rule-set rule-set-7a
```

In the following example, the `ds dd create` command creates a distributed device, and with the default rule-set:

```
Vplexcli:/> ds dd create --name TestDisDevice --devices  
/clusters/cluster-1/devices/TestDevCluster1, /clusters/cluster-2/devices/TestDevCluster2
```

Distributed-device 'TestDisDevice' is using rule-set 'cluster-2-detaches'.

- See also**
- ◆ [device attach-mirror on page 181](#)
  - ◆ [ds dd destroy on page 223](#)
  - ◆ [local-device create on page 279](#)

## ds dd declare-winner

Declares a winning cluster for a distributed-device that is in conflict after a link outage.

**Contexts** All contexts.

In /distributed-storage/distributed-device context, command is **declare-winner**.

In /distributed-storage context, command is **dd declare-winner**.

**Syntax**

```
ds dd declare-winner
  [-c|--cluster] context path
  [-d|--distributed-device] context path
  [-f|--force]
```

**Arguments** **Required arguments**

**[-c|--cluster] *context path* - \*** Specifies the winning cluster.

**[-d|--distributed-device] *context path* -** Specifies the distributed-device for which to declare a winning cluster.

**[-f|--force] -** Forces the 'declare-winner' command to be issued.

\* - argument is positional.

**Description** If the legs at two or more clusters are in conflict, use the **ds dd declare-winner** command to declare a winning cluster for a specified distributed device.

### Example

```
Vplexcli:/distributed-storage/distributed-devices> ds dd declare-winner --distributed-device
DDtest_4 --cluster cluster-2 --force
```

**See also** ♦ [ds dd create on page 219](#)

## ds dd destroy

Destroys the specified distributed-device(s).

**Contexts** All contexts.

**Syntax**

```
ds dd destroy
  [-d|--distributed-devices] context path, context path...
  [-f|--force]
```

**Arguments** **Required arguments**  
[-d|--distributed-device] *context path, context path,...* - \* List of one or more distributed-device(s) to destroy.

[-f|--force] - Forces the distributed-device(s) to be destroyed.

\* - argument is positional.

**Description** In order to be destroyed, the target distributed device must not host virtual volumes.

**Example**

```
Vplexcli:/distributed-storage/distributed-devices> ds dd destroy
/distributed-storage/distributed-devices/TestDisDevice
```

```
WARNING: The following items will be destroyed:
```

```
Context
```

```
-----
/distributed-storage/distributed-devices/TestDisDevice
```

```
Do you wish to proceed? (Yes/No) yes
```

**See also** ♦ [ds dd create on page 219](#)

## ds dd remove-all-rules

Removes all rules from all distributed devices.

**Contexts** All contexts.

**Syntax** ds dd remove-all-rules  
[-f|--force]

**Arguments** **Required arguments**

None.

**Optional arguments**

[-f|--force] - Force the operation to continue without confirmation.

**Description** From any context, removes all rules from all distributed devices.



### **WARNING**

*There is NO undo for this procedure.*

### **Example**

```
VPlexcli:/distributed-storage/distributed-devices/dd_23> remove-all-rules
```

```
All the rules in distributed-devices in the system will be removed. Continue? (Yes/No) yes
```

- See also**
- ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)

## ds dd set-log

Allocates/unallocates segments of a logging volume to a distributed device or a component of a distributed device.

**Contexts** All contexts.

**Syntax**

```
ds dd set-log
  [-d|--distributed devices] context path, context path...
  [-n|--cancel]
  [-c|--distributed-device-component] context path
  [-l|--logging-volumes] context path, context path...
```

**Arguments** **Required arguments**

**[-d|--distributed-devices] context path, context path...** - One or more distributed-device for which segments of the specified logging volume are allocated/unallocated.

All components of the distributed-device are included.

or:

**[-c|--distributed-device-component] context path** - The distributed-device component for which segments of the specified logging volume are allocated/unallocated.

and:

**[-l|--logging-volume] context path, context path...** - One or more logging-volume where the logging volume segments are allocated/unallocated. Target volume must be created as logging-volume. See [logging-volume create on page 293](#).

If not specified, the VPLEX configuration automatically allocates a logging volume for each cluster.

**Optional arguments**

**[-n|--cancel]** - Cancel/unallocate the log setting for the specified component of a distributed-device or all the components of the specified distributed-device. See WARNINGS below.

**Description** Logging volumes keep track of 4 k byte blocks written during an inter-cluster link failure. When the link recovers, VPLEX uses the information in logging volumes to synchronize the mirrors.



**WARNING**

*If no logging volume is allocated to a distributed device, a full rebuild of the device will occur when the inter-cluster link is restored after an outage.*



**WARNING**

*Do not change a device's logging volume unless the existing logging-volume is corrupted or unreachable or to move the logging volume to a new disk.*

Use the **ds dd set-log** command only to repair a corrupted logging volume or to transfer logging to a new disk.

Use the **--distributed-devices** argument to allocate/unallocate segments on the specified logging volume to the specified device.

Use the **--distributed-devices-component** argument to allocate/unallocate segments on the specified logging volume to the specified device component.

---

**Note:** Specify either distributed devices or distributed device components. Do not mix devices and components in the same command.

---

If the logging volume specified by the **--logging-volume** argument does not exist, it is created.

Use the **--cancel** argument to delete the log setting for a specified device or device component.



#### **WARNING**

---

*Use the **--cancel** argument very carefully.*

*No warning or confirmation message is displayed.*

*Removing the logging-volume for a device deletes the existing logging entries for that device. A FULL rebuild of the device will occur after a link failure and recovery.*

*Removing the logging volume for all distributed devices will remove all entries from the logging volume. In the event of a link failure and recovery, this results in a FULL rebuild of all distributed devices.*

---

**Example**    Allocate segments of a logging volume to a distributed device:

```
VPLEXcli:/distributed-storage/distributed-devices/TestDisDevice> ds dd set-log
--distributed-devices TestDisDevice --logging-volumes
/clusters/cluster-2/system-volumes/New-Log_Vol
```

Remove the logging volume for a distributed device:

```
VPLEXcli:/distributed-storage/distributed-devices/TestDisDevice> ds dd set-log
--distributed-devices TestDisDevice --cancel
```

**See also**    ♦ [logging-volume create on page 293](#)

## ds rule destroy

Destroys an existing rule.

**Contexts** All contexts.

In /distributed-storage context, command is **rule destroy**.

**Syntax** `ds rule destroy  
[-r|--rule} rule]`

**Arguments** **Required arguments**  
`[-r|--rule] rule -`

**Description** A rule-set contains rules. Use the **ll** command in the rule-set context to display the rules in the rule-set.

**Example** Use the **ds rule destroy** command to destroy a rule in the rule set.

```
Vplexcli:/distributed-storage/rule-sets/TestRuleSet/rules> ll
RuleName RuleType Clusters ClusterCount Delay Relevant
-----
rule_1    island-size      2          5s         true
```

```
Vplexcli:/distributed-storage/rule-sets/TestRuleSet/rules> rule destroy rule_1
```

- See also**
- ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)

## ds rule island-containing

Adds a island-containing rule to an existing rule-set.

**Contexts** All contexts.

In /distributed-storage context, command is **rule island-containing**.

**Syntax**

```
ds rule island-containing
[-c|--clusters] context path,context path...
[-d|--delay] delay
[-r|rule-set] context path
```

**Arguments** **Required arguments**

**[-c|--clusters] context path, context path...** - \* Clusters to which this rule applies.

**[-d|--delay] delay** - \* Sets the delay after a link outage before the rule is applied. Values must a positive integer and end with one of the following units:

**s** - delay timer is seconds (default unit)

**min** - delay timer is seconds

**h** - delay timer is hours

**[-r|--rule-set]** - Rule-set to which this rule is added.

\* - argument is positional.

**Description** Describes when to resume I/O on all clusters in the island containing the specified cluster.

**Example** In the following example, the **rule island-containing** command creates a rule that dictates:

- ◆ VPlex waits for 10 seconds after a link failure and then:
- ◆ Resumes I/O to the island containing cluster-1,
- ◆ Detaches any other islands.

```
VPlexcli:/distributed-storage/rule-sets/TestRuleSet/rules> ds rule island-containing
--clusters cluster-1 --delay 10s
```

```
VPlexcli:/distributed-storage/rule-sets/TestRuleSet/rules> ll
RuleName RuleType Clusters ClusterCount Delay Relevant
-----
rule_1 island-containing cluster-1 - 10s true
```

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)

---

## ds rule island-size

Creates a new island-size rule for a rule-set.

**Note:** This command is not supported in the current release.

**Contexts** All contexts.

In /distributed-storage context, command is **rule island-size**.

**Syntax**

```
ds rule island-size
[-c|--count] count
[-d|--delay] delay
[-r|--rule-set
```

**Arguments** **Required arguments**

**[-c|--count] count** - \* Number of clusters.

**[-d|--delay] delay** - \* Sets the delay after a link outage before the rule is applied. Values must be a positive integer and end with one of the following units:

**s** - delay timer is seconds (default unit)

**min** - delay timer is seconds

**h** - delay timer is hours

**[-r|--rule-set]** - Rule-set to which this rule is added.

\* - argument is positional.

**Description** Describes when to resume I/O on all clusters in the island containing the specified number of clusters.

Not supported in the current release.

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)

## ds rule-set copy

Copy an existing rule-set.

**Contexts** All contexts.

In /distributed-storage/rule-sets context, command is **copy**.

In /distributed-storage context, command is **rule-set copy**.

**Syntax**

```
ds rule-set copy
  [-s|--source] rule-set
  [-d|--destination] rule-set name
```

**Arguments** **Required arguments**

**[-s|--source] rule-set** - \* Source rule-set.

**[-d|--destination] new-rule-set** - The destination rule-set name.

\* - argument is positional.

**Description** Copies an existing rule-set and assigns the specified name to the copy.

### Example

```
Vplexcli:/distributed-storage/rule-sets> ll
```

```
Name          PotentialConflict  UsedBy
```

```
-----
TestRuleSet    false
```

```
Vplexcli:/distributed-storage/rule-sets> rule-set copy --source TestRuleSet --destination
```

```
CopyOfTest
```

```
Vplexcli:/distributed-storage/rule-sets> ll
```

```
Name          PotentialConflict  UsedBy
```

```
-----
CopyOfTest     false
```

```
TestRuleSet    false
```

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)

## ds rule-set create

Creates a new rule-set with the given name and encompassing clusters.

**Contexts** All contexts.

In /distributed-storage/rule-sets context, command is **create**.

In /distributed-storage context, command is **rule-set create**.

**Syntax** `ds rule-set create  
[-n|--name] rule-set`

**Arguments** **Required arguments**  
`[-n|--name] rule-set` - Name of the new rule-set.

**Examples** Create a rule-set:

```
Vplexcli: /> ds rule-set create --name TestRuleSet
```

| Name        | PotentialConflict | UsedBy |
|-------------|-------------------|--------|
| TestRuleSet | false             |        |

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)
  - ◆ [ds rule-set what-if on page 233](#)
  - ◆ [set on page 444](#)

## ds rule-set destroy

Destroys an existing rule-set.

**Contexts** All contexts.

In /distributed-storage/rule-sets context, command is **destroy**.

In /distributed-storage context, command is **rule-set destroy**.

**Syntax**

```
ds rule-set destroy
[-r|--rule-set] rule-set
```

**Arguments** **Required arguments**  
[-r|--rule-set] *rule-set* - Name of the rule-set to destroy.

**Description** Deletes the specified rule-set. The specified rule-set can be empty or can contain rules. Before deleting a rule-set, use the **set** command to detach the rule-set from any virtual volumes associated with the rule-set.

**Example** Delete a rule-set:

```
Vplexcli:/distributed-storage/rule-sets/NewRuleSet> ds rule-set destroy NewRuleSet
Context '/distributed-storage/rule-sets/NewRuleSet' has been removed.
```

In the following example:

- ◆ The **ll** command displays to which devices the rule-set is attached.
- ◆ The **set rule-set name ""** command detaches the rule set from a device.
- ◆ The **ds rule-set destroy** command deletes the rule-set.

```
Vplexcli:/distributed-storage/rule-sets/TestRuleSet> ll
```

Attributes:

| Name               | Value                    |
|--------------------|--------------------------|
| key                | ruleset_5537985253109250 |
| potential-conflict | false                    |
| used-by            | dd_00                    |

```
Vplexcli:/distributed-storage/rule-sets/TestRuleSet> cd
//distributed-storage/distributed-devices/dd_00
```

```
Vplexcli:/distributed-storage/distributed-devices/dd_00>set rule-set-name ""
Removing the rule-set from device 'dd_00' could result in data being unavailable during a WAN
link outage. Do you wish to proceed ? (Yes/No) yes
```

```
Vplexcli:/distributed-storage/distributed-devices/dd_00>ds rule-set destroy TestRuleSet
```

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set what-if on page 233](#)
  - ◆ [set on page 444](#)

---

## ds rule-set what-if

Tests if/when I/O is resumed at individual clusters, according to the current rule-set.

**Contexts** All contexts.

In /distributed-storage/rule-sets context, command is **what-if**.

In /distributed-storage context, command is **rule-set what-if**.

**Syntax**

```
ds rule-set what-if
[-i|--islands] "cluster-1,cluster-2"
[-r|--rule-set] context path
```

**Arguments** **Required arguments**

**[-i|--islands] "cluster-1,cluster-2"** - List of islands, in quotes. The clusters are separated by commas, the islands by a space.

**[-r|--rule-set] context path** - Context path of the rule-set used in the what-if scenario.

**Description** Only two clusters and one island are supported for the current release.

**Example** Test a rule-set:

```
Vplexcli:/distributed-storage/rule-sets> ds rule-set what-if --islands "cluster-1,cluster-2"
--rule-set TestRuleSet
IO does not stop.
```

- See also**
- ◆ [ds dd remove-all-rules on page 224](#)
  - ◆ [ds rule destroy on page 227](#)
  - ◆ [ds rule island-containing on page 228](#)
  - ◆ [ds rule island-size on page 229](#)
  - ◆ [ds rule-set copy on page 230](#)
  - ◆ [ds rule-set create on page 231](#)
  - ◆ [ds rule-set destroy on page 232](#)

## ds summary

Display summary information about distributed-devices.

**Contexts** All contexts.

In /distributed-storage context, command is **summary**.

**Syntax** ds summary

**Description** Displays summarized information for all distributed-devices.

Displays more detailed information for any device with a health-state or operational-status other than 'ok', and a service-status other than 'running'.

Displays devices per cluster, and calculates total and free capacity.

Use the **--verbose** argument to display additional information about unhealthy volumes in each consistency group.

**Example** Display summary information when no devices are unhealthy:

```
VPlexcli:/distributed-storage> ds summary
```

```
Distributed devices health summary:  
Total 25 devices, 0 unhealthy.
```

```
Cluster summary:  
  
Cluster cluster-2 : 25 distributed devices.  
Cluster cluster-1 : 25 distributed devices.
```

```
Capacity summary:  
  
0 devices have some free capacity.  
0B free capacity of 500G total capacity.
```

```
Distributed volume summary:  
Total 25 distributed devices in consistency groups, 0 unhealthy.  
Total 0 distributed devices not in consistency groups, 0 unhealthy
```

**Example** Display summary information when one or more devices are unhealthy:

```
VPlexcli:/> ds summary
```

```
Distributed Volumes (not in Consistency Groups) Unhealthy Summary:
```

| Device Name | Health State  | Operational Status | Service Status      |
|-------------|---------------|--------------------|---------------------|
| DR10        | major-failure | stressed           | cluster-unreachable |

```
Distributed volumes (in consistency groups) unhealthy summary:
```

| CG Name          | Cache Mode   | Number of Unhealthy Vols | Cluster                | Operational Status   | Status Details                                              |
|------------------|--------------|--------------------------|------------------------|----------------------|-------------------------------------------------------------|
| AA_ACW_Cluster12 | asynchronous | 9                        | cluster-1<br>cluster-2 | unknown<br>suspended | []<br>[cluster-departure,<br>restore-link-or-choose-winner] |
| AP_ACW_Cluster1  | asynchronous | 10                       | cluster-1<br>cluster-2 | unknown<br>suspended | []<br>[cluster-departure,<br>restore-link-or-choose-winner] |
| AP_ACW_Cluster2  | asynchronous | 5                        | cluster-1<br>cluster-2 | unknown<br>suspended | []<br>[cluster-departure,<br>restore-link-or-choose-winner] |

```
Distributed devices health summary:  
Total 25 devices, 25 unhealthy.
```

Cluster summary:

Cluster cluster-2 : 25 distributed devices.  
Cluster cluster-1 : 25 distributed devices.

Capacity summary:

0 devices have some free capacity.  
0B free capacity of 500G total capacity.

Distributed volume summary:

Total 24 distributed devices in consistency groups, 24 unhealthy.  
Total 1 distributed devices not in consistency groups, 1 unhealthy.

**Example** Use the `--verbose` argument to display detailed information about unhealthy volumes in each consistency group:

```
VPlexcli:/> ds summary --verbose
```

Distributed Volumes (not in Consistency Groups) Unhealthy Summary:

| Device Name | Health State  | Operational Status | Service Status      |
|-------------|---------------|--------------------|---------------------|
| DR10        | major-failure | stressed           | cluster-unreachable |

Distributed volumes (in consistency groups) unhealthy summary:

| CG Name          | Cache Mode   | Number of Unhealthy Vols | Cluster                | Operational Status   | Status Details                                                  |
|------------------|--------------|--------------------------|------------------------|----------------------|-----------------------------------------------------------------|
| AA_ACW_Cluster12 | asynchronous | 9                        | cluster-1<br>cluster-2 | unknown<br>suspended | [<br>[cluster-departure,<br>restore-link-or-choose-winner]<br>] |
| AP_ACW_Cluster1  | asynchronous | 10                       | cluster-1<br>cluster-2 | unknown<br>suspended | [<br>[cluster-departure,<br>restore-link-or-choose-winner]<br>] |
| AP_ACW_Cluster2  | asynchronous | 5                        | cluster-1<br>cluster-2 | unknown<br>suspended | [<br>[cluster-departure,<br>restore-link-or-choose-winner]<br>] |

Distributed volumes (in consistency groups) unhealthy details:

| CG Name          | Unhealthy Vols                                                                                                       |
|------------------|----------------------------------------------------------------------------------------------------------------------|
| AA_ACW_Cluster12 | ['DR11_vol', 'DR12_vol', 'DR13_vol', 'DR14_vol', 'DR15_vol', 'DR16_vol', 'DR17_vol', 'DR18_vol', 'DR19_vol']         |
| AP_ACW_Cluster1  | ['DR20_vol', 'DR21_vol', 'DR22_vol', 'DR23_vol', 'DR24_vol', 'DR25_vol', 'DR6_vol', 'DR7_vol', 'DR8_vol', 'DR9_vol'] |
| AP_ACW_Cluster2  | ['DRa_12_vol', 'DRb_12_vol', 'DRc_12_vol', 'DRd_12_vol', 'DRe_12_vol']                                               |

Distributed devices health summary:

Total 25 devices, 25 unhealthy.

Cluster summary:

Cluster cluster-2 : 25 distributed devices.  
Cluster cluster-1 : 25 distributed devices.

Capacity summary:

0 devices have some free capacity.  
0B free capacity of 500G total capacity.

Distributed volume summary:

Total 24 distributed devices in consistency groups, 24 unhealthy.  
Total 1 distributed devices not in consistency groups, 1 unhealthy.

Cluster cluster-1 : 25 distributed devices.

Table 14 ds summary field descriptions (1 of 2)

| Field                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Distributed Volumes (not in Consistency Groups) Unhealthy Summary:</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Device Name                                                               | Name of the device.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Health State                                                              | <p><b>major failure</b> - One or more children of the distributed device is out-of-date and will never rebuild, possibly because they are dead or unavailable.</p> <p><b>minor failure</b> - Either one or more children of the distributed device is out-of-date and will rebuild, or the Logging Volume for the distributed device is unhealthy.</p> <p><b>non-recoverable error</b> - VPLEX cannot determine the distributed device's Health state.</p> <p><b>ok</b> - The distributed device is functioning normally.</p> <p><b>unknown</b> - VPLEX cannot determine the device's health state, or the state is invalid.</p>                                                                                                                                                                                                                                                                                                                                                                              |
| Operational Status                                                        | <p><b>degraded</b> - The distributed device may have one or more out-of-date children that will eventually rebuild.</p> <p><b>error</b> - One or more components of the distributed device is hardware-dead.</p> <p><b>ok</b> - The distributed device is functioning normally.</p> <p><b>starting</b> - The distributed device is not yet ready.</p> <p><b>stressed</b> - One or more children of the distributed device is out-of-date and will never rebuild.</p> <p><b>unknown</b> - VPLEX cannot determine the distributed device's Operational state, or the state is invalid.</p>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Service Status                                                            | <p><b>cluster unreachable</b> - VPLEX cannot reach the cluster; the status is unknown.</p> <p><b>need resume</b> - The other cluster detached the distributed device while it was unreachable. The distributed device needs to be manually resumed for I/O to resume at this cluster.</p> <p><b>need winner</b> - All clusters are reachable again, but both clusters had detached this distributed device and resumed I/O. You must pick a winner cluster whose data will overwrite the other cluster's data for this distributed device.</p> <p><b>potential conflict</b> - The clusters have detached each other resulting in a potential for detach conflict.</p> <p><b>running</b> - The distributed device is accepting I/O.</p> <p><b>suspended</b> - The distributed device is not accepting new I/O; pending I/O requests are frozen.</p> <p><b>winner-running</b> - This cluster detached the distributed device while the other cluster was unreachable, and is now sending I/O to the device.</p> |
| Cluster Summary                                                           | Number of distributed devices on each cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Capacity Summary                                                          | Number of devices with free capacity, amount of free capacity for the cluster, and total capacity for all clusters.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Distributed volumes (in consistency groups) unhealthy summary:</b>     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| CG Name                                                                   | Name of the consistency group of which the unhealthy device is a member.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Cache Mode                                                                | <p>Cache mode of the consistency group.</p> <p><b>Synchronous</b> - Supported on VPLEX Local and VPLEX Metro configurations where clusters are separated by up to 5 ms of latency. In synchronous cache mode, writes to the back-end storage volumes are not acknowledged to the host until the back-end storage volumes acknowledge the write.</p> <p><b>Asynchronous</b> - Supported on VPLEX Geo configurations where clusters separated by up to 50 ms of latency. In asynchronous cache mode, host writes to a distributed volume are acknowledged back to the host after the data is protected in the cache of another director in the local cluster.</p>                                                                                                                                                                                                                                                                                                                                               |
| Number of unhealthy volumes                                               | Number of unhealthy volumes in the consistency group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

Table 14 ds summary field descriptions (2 of 2)

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster            | <p>Visibility of the consistency group.</p> <p><b>cluster-1</b> - consistency group is visible only at cluster-1.</p> <p><b>cluster-2</b>- consistency group is visible only at cluster-2.</p> <p><b>cluster-1, cluster-2</b> - consistency group is visible at both clusters.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Operational Status | <p>Current status for this consistency group with respect to each cluster on which it is visible.</p> <p><b>ok</b> - I/O can be serviced on the volumes in the consistency group.</p> <p><b>suspended</b> - I/O is suspended for the volumes in the consistency group. The reasons are described in the "operational status: details".</p> <p><b>degraded</b> - I/O is continuing, but there are other problems described in "operational status: details":</p> <p><b>unknown</b> - The status is unknown, likely because of lost management connectivity.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Status Details     | <p>If operational status is "ok" this field is empty: "[ ]". Otherwise, it displays additional information, which may be any of the following:</p> <p><b>temporarily-synchronous</b> - This asynchronous consistency group has temporarily switched to 'synchronous' cache mode because of a director failure.</p> <p><b>requires-resolve-conflicting-detach</b> - After the inter-cluster link is restored, two clusters have discovered that they have detached one another and resumed I/O independently. The clusters are continuing to service I/O on their independent versions of the data. The consistency-group resolve-conflicting-detach command must be used to make the view of data consistent again at the clusters.</p> <p><b>rebuilding-across-clusters</b> - One or more distributed member volumes is being rebuilt. At least one volume in the group is out of date at that cluster and is re-syncing. If the link goes out at this time the entire group is suspended. Use the rebuild status command to display which volume is out of date at which cluster.</p> <p><b>rebuilding-within-cluster</b> - One or more local rebuilds is in progress at this cluster.</p> <p><b>data-safe-failure</b> - A single director has failed. The volumes are still crash-consistent, and will remain so, unless a second failure occurs before the first is recovered.</p> <p><b>requires-resume-after-data-loss-failure</b> - There have been at least two concurrent failures, and data has been lost. For example, a director fails shortly after the inter-cluster link fails, or when two directors fail at almost the same time. Use the consistency-group resume-after-data-loss-failure command to select a winning cluster and allow I/O to resume.</p> <p><b>cluster-departure</b> - Not all the visible clusters are in communication.</p> <p><b>requires-resume-after-rollback</b> - A cluster has detached its peer cluster and rolled back the view of data, but is awaiting the consistency-group resume-after-rollback command before resuming I/O.</p> <p>Displayed:</p> <ul style="list-style-type: none"> <li>• For an asynchronous consistency group where both clusters have been writing, that is where both clusters are active.</li> <li>• At the winning side when a detach rule fires, or shortly after the consistency-group choose-winner command picks a winning cluster.</li> </ul> <p><b>requires-resume-at-loser</b> - Displayed on the losing side when the inter-cluster link heals after an outage. After the inter-cluster link is restored, the losing cluster discovers that its peer was declared the winner and resumed I/O. Use the consistency-group resume-at-loser command to make the view of data consistent with the winner, and to resume I/O at the loser.</p> <p><b>restore-link-or-choose-winner</b> - I/O is suspended at all clusters because of a cluster departure, and cannot automatically resume. This can happen if:</p> <ul style="list-style-type: none"> <li>• There is no detach-rule</li> <li>• If the detach-rule is 'no-automatic-winner', or</li> <li>• If the detach-rule cannot fire because its conditions are not met.</li> </ul> <p>For example, if more than one cluster is active at the time of an inter-cluster link outage, the 'active-cluster-wins' rule cannot take effect. When this detail is present, I/O will not resume until either the inter-cluster link is restored, or the user intervenes to select a winning cluster with the consistency-group choose-winner command.</p> <p><b>unhealthy-devices</b> - I/O has stopped in this consistency group because one or more volumes is unhealthy and cannot perform I/O.</p> <p><b>will-rollback-on-link-down</b> - If there were a link-down now, the winning cluster would have to roll back the view of data in order to resume I/O.</p> |

- See also**
- ◆ [export port summary on page 248](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [extent summary on page 270](#)
  - ◆ [local-device summary on page 283](#)
  - ◆ [storage-volume summary on page 469](#)
  - ◆ [virtual-volume summary on page 510](#)

## event-test

Verifies that the management server can receive events from a director.

**Contexts** All contexts.

**Syntax**

```
event-test
[-n|--directors] context path,context path...
[-c|--clusters] context path,context path...
[-l|--level] [emergency|alert|critical|error|warning|
             notice|info|debug]
[-o|--component] component
[-m|--message] "message"
```

**Arguments** **Required arguments**

**[-l|--level] level** - Level of the event. Must be one of the following:

**emergency** - System is unusable.

**alert** - Immediate action is required.

**critical** - Critical condition detected.

**error** - Significant error condition detected.

**warning** - Warning condition is detected.

**notice** - Normal, but significant condition.

**info** - Information messages.

**debug** - Detailed event information used by EMC for debugging.

**Optional arguments**

**[-d|--directors] context path, context path...** - One or more directors from which to send the test event.

**[-c|--clusters] context path,context path...** - One or more clusters from which to send the test event. Events are sent from every director in the specified cluster(s).

**[-o|--component]** - Text to include in the component portion of the test message.

**[-m|--message] "message"** - Text of message to send in event test, enclosed in quotes. This text is written to the firmware log prefixed by "EVENT-TEST".

**Description** Tests the logging path from one or mores director to the management server.

Every component in the software that runs on the director logs messages to signify important events. Each logged message/event is transferred from the director to the management server and into the firmware log file.

Use this command to verify this logging path is ok. Specify what level of event should be generated. Optionally, specify the text to appear in the component portion of the test message.

Check the appropriate firmware log for the event created by this command.

**Example** In the following example:

- ◆ The **event test** command creates an alert event for the specified director.
- ◆ The **exit** command exits the CLI.
- ◆ The **tail** command displays the firmware log for the director.

```
Vplexcli: /> event-test --director director-2-1-A --level alert --message "Test Alert"
```

```
Vplexcli:/> exit  
Connection closed by foreign host.
```

```
service@ManagementServer:~> tail /var/log/Vplex/cli/firmware.log_20100903092147  
128.221.253.67/xmmg/log:5988:W/"2343":324:<1>2010/09/07 15:18:15.82: test/1 EVENT-TEST: Test  
Alert
```

---

## exec

Executes an external program.

**Contexts** All contexts.

**Syntax** `exec command name word*`

**Description** The program can be executed with zero or more arguments.

---

**Note:** The correct syntax for program names and arguments depends on the host system.

---

**Example** To display the date and time on Director-1-1-A:

```
Vplexcli:/> exec ssh 128.221.253.35 date  
Tue Sep 21 14:32:52 UTC 2010
```

## exit

Exits the shell.

**Contexts** All contexts.

**Syntax**

```
exit
[-e|--exit-code] exit-code
[-s|--shutdown]
```

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-e|--exit-code] *exit code*** - Returns the specified value when the shell exits. If no exit code is specified, then 0 is returned.

**[-s|--shutdown]** - When running in server mode, shuts down the shell instead of closing the socket. No effect if not running in server mode.

**Description** If the shell is not embedded in another application, the shell process will stop.

**Example**

```
VPlexcli:/> exit
Connection closed by foreign host.
service@ManagementServer:~>
```

## export initiator-port discovery

Discovers initiator ports on the front-end fabric.

**Contexts** Cluster context and below.

In `/clusters/cluster/exports` context, command is **initiator-port discovery**.

In `/clusters/cluster/exports/initiator-ports` context, command is **discovery**.

**Syntax**

```
export initiator-port discovery
[-t|--timeout] 1 - 3600
[-w|--wait] 1 - 3600
[-c|--cluster] context path
```

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-t|--timeout] seconds` - The maximum number of seconds to wait for the front-end fabric discovery operation to complete.

Default: 300.

Range: 1- 3600.

`[-w|--wait]` - The maximum number of seconds to wait for a response from the fabric discovery.

Default: 10.

Range: 1- 3600.

`[-c|--cluster] context path` - Discover initiator ports on the specified cluster.

**Description** Initiator discovery finds unregistered initiator-ports on the front-end fabric and determines the associations between the initiator ports and the target ports.

Use the **ll** command in initiator-ports context to display the same information for small configurations (where timeout does not occur)

Use the **export initiator-port discovery** command for large configurations in which **ls** command may encounter timeout limits.

**Example** Discover initiator ports on another cluster:

```
Vplexcli:/clusters/cluster-1/exports/initiator-ports> discovery --cluster cluster-2
Name          port-wwn          node-wwn          type          Target Port Names
-----
LicoJ013_hba1 0x10000000c97b1f3d 0x10000000c97b1f3d sun-vcs
LicoJ009_hba1 0x10000000c992c841 0x10000000c992c841 sun-vcs
LicoJ007_hba3 0x10000000c98a9dae 0x10000000c98a9dae sun-vcs
LicoJ011_hba2 0x10000000c992bf61 0x10000000c992bf61 sun-vcs
LicoJ010_hba1 0x10000000c992c84b 0x10000000c992c84b sun-vcs
P000000003CA000E6-A1-FC00,
P000000003CA001CB-A1-FC00,
P000000003CB000E6-B1-FC00,
P000000003CB001CB-B1-FC00
.
.
.
```

**See also** ♦ [export initiator-port register on page 244](#)

## export initiator-port register

Registers an initiator-port and associates one WWN pair with it.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **initiator-port register**.

In `/clusters/cluster/exports/initiator-ports` context, command is **register**.

**Syntax**

```
export initiator-port register
  [-c|--cluster] context path
  [-t|--type] {hpux|sun-vcs|aix|recoverpoint|default}
  [-i|--initiator-port] initiator-port
  [-p|--port] port
```

**Arguments** **Required arguments**

**[-i|--initiator-port] *initiator-port*** - \* Name to assign to the registered port. Command fails if the specified name is already in use.

**[-p|--port] *port*** - \* The portWWN and nodeWWN where each WWN is either '0x' followed by one or more hexadecimal digits or an abbreviation, in the following format:

```
<string>:<number>[,<number>]
```

For example:

```
0xd1342a|0xd1342b
```

```
hyy1:194e,4|hyy1:194e
```

```
0xd1342a
```

```
hyy1:194e,4
```

The nodeWWN can be optional.

\* - argument is positional.

**Optional arguments**

**[-c|--cluster] *context path*** - Cluster on which the initiator port is registered.

**[-t|--type] {hpux|sun-vcs|aix|recoverpoint|default}** - Type of initiator port.

**hpux** - Hewlett Packard UX

**sun-vcs** - Sun Solaris

**aix** - IBM AIX

**recoverpoint** - EMC RecoverPoint

**default** - If no type is specified

**Description**

An initiator must be registered in order to see any exported storage. The initiator must also be able to communicate with the front-end ports. Registering attaches a meaningful name to the port's WWN and allows easier auditing of export settings.

Initiators can include port WWNs of HBAs that are not currently connected to the fabric. This allows system configuration even when the HBA is offline.

Use the **ll** command in `/engines/engine/directors/director/hardware /ports/port` context to display portWWNs and nodeWWNs.

**Example** Register a host initiator port:

- ◆ WWN - 0x10000000c97b0726
- ◆ Name the registered port iE\_209\_hba0
- ◆ FE port pair with node name 0x200100601610428f:

```
Vplexcli:/clusters/cluster-1/exports/initiator-ports> register --initiator-port iE_209_hba0  
--port 0x10000000c97b0726|0x200100601610428f
```

```
Vplexcli:/clusters/cluster-1/exports/initiator-ports> ll  
Name          port-wwn          node-wwn          type          Target Port Names  
-----  
iE_209_hba0   0x10000000c97b0726 0x200100601610428f default        P000000601610428F-A0-FC00,  
P000000601610672E-B0-FC00
```

**Note:** The FE node name represents one port from each director. In the example above, the FE pair is ports A0 and B0. For two engines there would be four ports, and with four engines there would be eight ports.

- See also**
- ◆ [export initiator-port discovery on page 243](#)
  - ◆ [export initiator-port unregister on page 247](#)
  - ◆ [export target-port renamewwns on page 266](#)

## export initiator-port register-host

Creates a view, and registers each port WWN /name pair as an initiator port in that view.

**Contexts** All contexts.

In /clusters/*cluster*/exports context, command is **initiator-port register-host**.

In /clusters/*cluster*/exports/initiator-ports context, command is **register-host**.

**Syntax**

```
export initiator-port register-host
  [-p|--ports] port,port...
  [-f|--file] file
  [-c|--cluster] cluster context
```

**Arguments** **Required arguments**

**[-f|--file] *file*** - \* The host declaration file path name.

**Optional arguments**

**[-c|--cluster] *cluster context*** - \* The cluster at which to create the view.

**[-p|--ports] *port,port...*** - List of port names. If omitted, all ports at the cluster will be used. Entries must be separated by commas.

\* - argument is positional.

**Description**

Reads host port WWNs (with optional node WWNs) and names from a host declaration file. Creates a view, registering each port WWN /name pair as an initiator port in that view.

The host description file contains one line for each port on the host in the following format:

```
<port WWN> [|node WWN>] <port name>
```

Hosts must be registered in order to be exported (added to a storage view).

Registering consists of naming the initiator and listing its ports WWN/GUID.

Each port of a server's HBA/HCA must be registered as a separate initiator.

**Example**

**See also**

- ◆ [export initiator-port discovery on page 243](#)
- ◆ [export initiator-port unregister on page 247](#)

---

## export initiator-port unregister

Unregisters the specified initiator-port(s).

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **initiator-port unregister**.

In `/clusters/cluster/exports/initiator-ports` context, command is **unregister**.

**Syntax**

```
export initiator-port unregister
    [-f|--force]
    [-i|--initiator-port] context path, context path...
```

**Arguments** **Required arguments**

`[-i|--initiator-port] context path, context path...` - \* One or more initiator ports to remove. Entries must be separated by commas.

**Optional arguments**

`[-f|--force]` - Destroys the initiator-ports even if they are in use.

\* - argument is positional.

**Example**

```
Vplexcli:> export initiator-port unregister -i win2k3_105_port1
```

**See also**

- [export initiator-port register on page 244](#)

## export port summary

Displays a summary of unhealthy exported ports.

**Contexts** All contexts.

In /clusters/*cluster*/exports context, command is **port summary**.

In /clusters/*cluster*/exports/ports context, command is **summary**.

**Syntax** export port summary  
[-c|--clusters] cluster,cluster

**Arguments** **Required arguments**

None.

**Optional arguments**

[-c|--clusters] *cluster,cluster...* - Display unhealthy ports for only the specified cluster(s).

**Description** Prints a summary of the unhealthy ports, the ports at each cluster, and the views which contain it.

In the root context, displays information for all clusters.

In /cluster context or below, displays information for only the current cluster.

**Example** Display port health for a specified cluster:

```
Vplexcli:/> export port summary --clusters cluster-1
```

```
Port health summary(cluster-1):
```

| port name                 | export status | view summary       |       |                 |
|---------------------------|---------------|--------------------|-------|-----------------|
| P000000003CA00147-A0-FC01 | suspended     | no unhealthy views |       |                 |
| P000000003CA00147-A0-FC03 | suspended     | no unhealthy views |       |                 |
| P000000003CA00147-A1-FC01 | suspended     | no unhealthy views |       |                 |
| .                         |               |                    |       |                 |
| .                         |               |                    |       |                 |
| .                         |               |                    |       |                 |
| port name                 | health state  | enabled            | views | virtual-volumes |
| P000000003CA00147-A0-FC00 | healthy       | true               | 1     | 1               |
| P000000003CA00147-A0-FC02 | healthy       | true               | 2     | 28              |
| P000000003CA00147-A1-FC00 | healthy       | true               | 1     | 1               |
| .                         |               |                    |       |                 |
| .                         |               |                    |       |                 |
| .                         |               |                    |       |                 |
| P000000003CB00147-B0-FC03 | unhealthy     | true               | 0     | 0               |
| P000000003CB00147-B1-FC01 | unhealthy     | true               | 0     | 0               |
| P000000003CB00147-B1-FC03 | unhealthy     | true               | 0     | 0               |

Total 16 ports, 8 unhealthy.

Display port health using verbose argument:

```
Vplexcli:/clusters/cluster-1/exports/ports> summary --verbose
```

```
Port health summary(cluster-1):
```

| port name                 | export status | view summary       |
|---------------------------|---------------|--------------------|
| P000000003CA00147-A0-FC01 | suspended     | no unhealthy views |
| P000000003CA00147-A0-FC03 | suspended     | no unhealthy views |

```

P000000003CA00147-A1-FC01  suspended      no unhealthy views
.
.
.
port name                    health state  enabled  view summary
-----
P000000003CA00147-A0-FC00  healthy      true     view      virtual-volumes
                               LicoJ013      1
P000000003CA00147-A0-FC02  healthy      true     view      virtual-volumes
                               LicoJ013      1
                               LicoJ009      27
.
.
.
P000000003CB00147-B0-FC03  unhealthy    true     no views
P000000003CB00147-B1-FC01  unhealthy    true     no views
P000000003CB00147-B1-FC03  unhealthy    true     no views

```

Total 16 ports, 8 unhealthy.

- See also**
- ◆ [ds summary on page 234](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [extent summary on page 270](#)
  - ◆ [local-device summary on page 283](#)
  - ◆ [storage-volume summary on page 469](#)
  - ◆ [virtual-volume summary on page 510](#)

## export storage-view addinitiatorport

Adds the specified initiator port(s) to a storage view.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **storage-view addinitiatorport**.

In `/clusters/cluster/exports/storage-views` context, command is **addinitiatorport**.

**Syntax** `export storage-view addinitiatorport  
[-v|--view] context path  
[-i|--initiator-ports] context path,context path...`

**Arguments** **Required arguments**

`[-i|--initiator-ports] context path,context path ...` - \* List of one or more initiator ports to add to the view. Entries must be separated by commas.

**Optional arguments**

`[-v|--view] context path` - View to which to add the specified initiator port(s).

\* - argument is positional.

**Description** Select ports from two different directors so as to maximize redundancy.

**Example** Add the initiator `iE_209_hba0` to the view named `Dell_209_view`:

```
VPLEXcli:/clusters/cluster-1/exports> storage-view addinitiatorport --view Dell_209_view  
--initiator-ports iE_209_hba0
```

**See also**

- ◆ [export storage-view create on page 255](#)
- ◆ [export storage-view removeinitiatorport on page 260](#)

---

## export storage-view addport

Adds the specified port(s) to the storage view.

**Contexts** All contexts.

In clusters/*cluster*/exports/storage-views/*storage-view* context, command is **addport**.

**Syntax**

```
export storage-view addport
[-v|--view] context path
[-p|--ports] context path,context path...
```

**Arguments**

**Required arguments**  
[-p|--ports] *context path,context path* ... - \* List of one or more ports to be added to the view. Entries must be separated by commas.

**Optional arguments**  
[-v|--view] *context path* - Storage view to which to add the specified port(s).

\* - argument is positional.

**Description** Use the `ll /clusters/cluster/exports/ports` command to display ports on the cluster.

**Example**

```
Vplexcli:/clusters/cluster-1/exports/storage-views/TestStorageView>
export storage-view addport --ports P000000003CB00147-B0-FC03
```

**See also**

- ◆ [export storage-view create on page 255](#)
- ◆ [export storage-view removeport on page 261](#)

## export storage-view addvirtualvolume

Adds a virtual volume to a storage view.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **storage-view addvirtualvolume**.

In `/clusters/cluster/exports/storage-views` context, command is **addvirtualvolume**.

**Syntax**

```
export storage-view addvirtualvolume
  [-v|--view] context path
  [-f|--force]
  [-o|--virtual-volumes] virtual-volume, virtual-volume...
```

**Arguments** **Required arguments**

**[-o|--virtual-volumes] virtual volume, virtual volume ...** - \* List of one or more virtual-volumes or LUN-virtual-volume pairs. Entries must be separated by commas.

LUN-virtual-volume pairs must be enclosed in parentheses (). Virtual-volumes and LUN-virtual-volume pairs can be typed on the same command line.

When only virtual-volumes are specified, the next available LUN is automatically assigned by VPLEX.

**Optional arguments**

**[-v|--view] context path>** - View to which to add the specified virtual volume(s).

**[-f|--force]** - Force the virtual-volumes to be added to the view even if they are already in use, if they are already assigned to another view, or if there are problems determining the view's state. Virtual-volumes that already have a LUN in the view will be re-mapped to the newly-specified LUN.

\* - argument is positional.

**Description** Add the specified virtual volume to the specified storage view. Optionally, specify the LUN to assign to the virtual volume. Virtual volumes must be in a storage view in order to be accessible to hosts.

When virtual-volumes are added using only volume names, the next available LUN number is automatically assigned.

Virtual-volumes and LUN-virtual-volume pairs can be specified in the same command line. For example:

```
r0_1_101_vol, (2, r0_1_102_vol), r0_1_103_vol
```

To modify the LUN assigned to a virtual volume, specify a virtual volume that is already added to the storage view and provide a new LUN.

**Example** Add a virtual-volume `Symm1254_7BF_1_vol` to the storage view `E_209_view`:

```
Vplexcli:/clusters/cluster-1/exports> storage-view addvirtualvolume --view E_209_view
--virtual-volumes Symm1254_7BF_1_vol
```

**Example** Modify the LUN assigned to a virtual volume already added to a view:

- ◆ The **ll** command in storage view context displays the LUN (0) assigned to a storage volume.
- ◆ The **export storage-view addvirtualvolume (LUN,Virtual-volume) --force** command assigns a new LUN to the virtual volume.

- ◆ The **ll** command in storage view context displays the new LUN assigned to a storage volume:

```
Vplexcli:/clusters/cluster-1/exports/storage-views/TestStorageView> ll
Name                               Value
-----
controller-tag                     -
initiators                         []
operational-status                 stopped
port-name-enabled-status           [P000000003CA00147-A1-FC01,true,suspended,
P000000003CB00147-B0-FC01,true,suspended]
ports                               [P000000003CA00147-A1-FC01, P000000003CB00147-B0-FC01]
virtual-volumes                    [(0,TestDisDevice_vol,VPD83T3:600014400000010a0014760d64cb325,16G)]
```

```
Vplexcli:/clusters/cluster-1/exports/storage-views/TestStorageView> export storage-view
addvirtualvolume (5,TestDisDevice_vol) --force
WARNING: Volume 'TestDisDevice_vol' already has LUN 0 in this view; remapping to LUN 5.
```

```
Vplexcli:/clusters/cluster-1/exports/storage-views/TestStorageView> ll
Name                               Value
-----
controller-tag                     -
initiators                         []
operational-status                 stopped
port-name-enabled-status           [P000000003CA00147-A1-FC01,true,suspended,
P000000003CB00147-B0-FC01,true,suspended]
ports                               [P000000003CA00147-A1-FC01, P000000003CB00147-B0-FC01]
virtual-volumes                    [(5,TestDisDevice_vol,VPD83T3:600014400000010a0014760d64cb325,16G)]
```

**Example** Add a synchronous virtual volume to a view using the **--force** option from the root context:

```
Vplexcli:/> export storage-view addvirtualvolume --view
/clusters/Saul1/exports/storage-views/TestStorageView --virtual-volumes
dr710_20_C1Win_0038_12_vol --force
```

Volume {1} is synchronous and on a non-local device. Applications using this volume may experience per I/O inter-cluster latency. If the applications are sensitive to this latency, they may experience data unavailability. Do you wish to proceed? (Yes/No)

- See also**
- ◆ [export storage-view checkconfig on page 254](#)
  - ◆ [export storage-view create on page 255](#)
  - ◆ [export storage-view removevirtualvolume on page 262](#)

## export storage-view checkconfig

Checks the configuration of the views.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **storage-view checkconfig**.

In `/clusters/cluster/exports/storage-views` context, command is **checkconfig**.

**Syntax** `export storage-view checkconfig`

**Description** Performs the following checks:

- ◆ Identifies any virtual volumes that are exported more than once.
- ◆ Identifies views that contain only a single port.
- ◆ Identifies views that are disabled.

**Example** To check all view configurations for all clusters from the CLI, type:

```
Vplexcli:/> export storage-view checkconfig
Checking cluster cluster-1:
No errors found for cluster cluster-1.
Checking cluster cluster-2:
No errors found for cluster cluster-2.
Volume dd_13_vol is exported multiple times:
  view: LicoJ009, lun: 14
  view: LicoJ010, lun: 14
Volume dd_16_vol is exported multiple times:
  view: LicoJ009, lun: 17
  view: LicoJ010, lun: 17
Volume dd_12_vol is exported multiple times:
  view: LicoJ009, lun: 13
  view: LicoJ010, lun: 13
Volume dd_19_vol is exported multiple times:
  view: LicoJ009, lun: 20
  view: LicoJ010, lun: 20
.
.
.
```

- See also**
- ◆ [export storage-view create on page 255](#)
  - ◆ [export storage-view find on page 257](#)
  - ◆ [export storage-view map on page 259](#)
  - ◆ [export storage-view show-powerpath-interfaces on page 263](#)

## export storage-view create

Creates a view with the given ports.

**Contexts** All contexts.

**Syntax**

```
export storage-view create
  [-c|--cluster] context path
  [-n|--name] name
  [-p|--ports] context path,context path...
```

**Arguments**

**Required arguments**

`[-n|--name] name` - \* Name of the new view. Must be unique throughout VPLEX.

`[-p|--ports] context path,context path...` - \* List of one or more ports to add to the view.

**Optional arguments**

`[-c|--cluster] context path` - The cluster on which to create the view.

\* - argument is positional.

**Description** A storage view is a logical grouping of front-end ports, registered initiators (hosts), and virtual volumes used to map and mask LUNs. Storage views are used to control host access to storage.

In order for hosts to access virtual volumes, the volumes must be in a storage view. A storage view consists of:

- ◆ One or more initiators. Initiators are added to a storage view using the “[export storage-view addinitiatorport](#)” command.
- ◆ One or more virtual-volumes. Virtual volumes are added to a storage view using the “[export storage-view addvirtualvolume](#)” command.
- ◆ One or more FE ports. Ports are added to a storage view using the “[export storage-view addport](#)” command.



### CAUTION

The name assigned to the storage view must be unique throughout the VPLEX. In VPLEX Metro or Geo configurations, the same name must not be assigned to a storage view on the peer cluster.

Use the `ll clusters/*/exports/storage-views` command to display the names of existing storage views before assigning a name.

**Example** Create a view named E\_209\_view for FE ports A0 and B0:

```
Vplexcli:/clusters/cluster-1/exports/storage-views> storage-view create --cluster
/clusters/cluster-1 --name E_209_View --ports
P000000601610428F-A0-FC00,P000000601610672E-B0-FC00
```

**See also**

- ◆ [export storage-view addport](#) on page 251
- ◆ [export storage-view addinitiatorport](#) on page 250
- ◆ [export storage-view addvirtualvolume](#) on page 252
- ◆ [export storage-view destroy](#) on page 256

## export storage-view destroy

Destroys the specified storage view.

**Contexts** All contexts.

**Syntax** `export storage-view destroy`  
`[-f|--force]`  
`[-v|--view] context path`

**Arguments** **Required arguments**  
`[-v|--view] context path ...` - \* Storage view to destroy.

**Optional arguments**  
`[-f|--force]` - Force the storage view to be destroyed even if it is in use.  
\* - argument is positional.

**Description** Destroys the specified storage view.

**Example** `Vplexcli:/> export storage-view destroy`  
`/clusters/cluster-1/exports/storage-views/TestStorageView`

**See also**

- ◆ [export storage-view create on page 255](#)
- ◆ [export storage-view removeinitiatorport on page 260](#)
- ◆ [export storage-view removeport on page 261](#)
- ◆ [export storage-view removevirtualvolume on page 262](#)

## export storage-view find

Displays export views for a specified volume, LUN, initiator, or cluster. Displays next available LUN number for all storage views.

**Contexts** Cluster/exports and below.

In `/clusters/cluster/exports` context, command is **storage-view find**.

In `/clusters/cluster/exports/storage-views` context, command is **find**.

**Syntax**

```
export storage-view find
  [-c|--cluster] cluster]
  [-v|--volume] volume
  [-l|--lun] LUN
  [-i|--initiator] initiator
  [-f|--free-lun]
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**[-c|--cluster] cluster** - Cluster to search for views.

**[-v|--volume] volume** - Find the views exporting the specified volume. Identify the volume by name, VPD83 identifier, or a name pattern with wildcards.

**[-l|--lun] LUN** - Find the views exporting the specified LUN number.

**[-i|--initiator-port] initiator** - Find the views including the specified initiator. May contain wildcards.

**[-f|--free-lun]** - Find the next free LUN number for all views.

**Description** This command is most useful for configurations with thousands of LUNs, and a large number of views and exported virtual volumes.

**Example** Find the next available LUN numbers on cluster 1:

```
Vplexcli:/clusters/cluster-1/exports/storage-views> find --cluster cluster-1 --free-lun
View LicoJ009 : next free LUN number is 27.
View LicoJ013 : next free LUN number is 1.
```

Find the views exporting the specified volume:

```
Vplexcli:/clusters/cluster-1/exports> export storage-view find --volume dd_03_vol
Views exporting volume dd_03_vol:
View LicoJ009 exports (4,dd_03_vol,VPD83T3:6000144000000010a0014760d64cb225,128G).
```

Find the views exported by initiators whose name starts with "Lico":

```
Vplexcli:/clusters/cluster-1/exports> export storage-view find --initiator Lico*
Views including initiator Lico*:
View LicoJ009.
View LicoJ013.
```

- See also**
- ◆ [export initiator-port discovery on page 243](#)
  - ◆ [export storage-view find-unmapped-volumes on page 258](#)
  - ◆ [export storage-view map on page 259](#)
  - ◆ [export storage-view summary on page 264](#)

## export storage-view find-unmapped-volumes

Displays unexported virtual volumes.

**Contexts** All contexts.

**Syntax** `export storage-view find-unmapped-volumes [-c|--cluster] cluster`

**Arguments** **Required arguments**  
`[-c|--cluster] cluster` - Cluster for which to display unexported storage volumes.

**Description** Displays unexported virtual volumes in the specified cluster.  
Displays the remote (on the other cluster) virtual volumes which are exported and running I/O.

**Example** Display unexported volumes for the specified clusters:

```
Vplexcli:/> export storage-view find-unmapped-volumes -c cluster-1
Name                               Operational Health Service Block Block Capacity Locality Supporting Device Cache Mode Expandable
----- Status State Status Count Size -----
device_Symm1690t1_00A8_1_vol        ok      ok running 528000 4K 2.01G local device_Symm1690t1_00A8_1 synchronous true
dr_1_vol                            ok      ok running 528000 4K 2.01G distributed dr_1 synchronous false

Vplexcli:/> export storage-view find-unmapped-volumes -c cluster-2
Name                               Operational Health Service Block Block Capacity Locality Supporting Device Cache Mode Expandable
----- Status State Status Count Size -----
device_Symm0381t1_0850_1_vol        ok      ok running 528000 4K 2.01G local device_Symm0381t1_0850_1 synchronous true
device_Symm0381t1_0851_1_vol        ok      ok running 528000 4K 2.01G local device_Symm0381t1_0851_1 synchronous true
.
.
.
device_Symm1690t1_00AB_1_vol        ok      ok running 528000 4K 2.01G remote device_Symm1690t1_00AB_1 synchronous false
device_Symm1690t1_00AC_1_vol        ok      ok running 528000 4K 2.01G remote device_Symm1690t1_00AC_1 synchronous false
```

**See also**

- ◆ [export storage-view addvirtualvolume on page 252](#)
- ◆ [export storage-view removevirtualvolume on page 262](#)

## export storage-view map

Displays only storage volumes with an I/O status other than 'Alive'.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **view map**.

In `/clusters/cluster/exports/storage-views` context, command is **map**.

**Syntax**

```
export storage-view map
  [-v|--views] view,view...
  [-f|--file] filename
```

**Arguments** **Required arguments**

`[-v|--views] view,view...` - \* List of one or more storage views to map. Entries must be separated by commas. May contain wildcards.

**Optional arguments**

`[-f|--file] file` - Name of the file to which to send the output. If no file is specified, output is to the console screen.

\* argument is positional.

**Example** Display unhealthy storage volumes for a specified storage view:

```
Vplexcli:/> export storage-view map LicoJ013
VPD83T3:6000144000000010a0014760d64cb32c dev_sym1723_1FC_vol
```

Display unhealthy storage volumes for all storage views:

```
Vplexcli:/> export storage-view map --views **
VPD83T3:6000144000000010a0014760d64ca44c base0_vol
VPD83T3:6000144000000010a0014760d64cb21f dd_00_vol
.
.
.
```

Display unhealthy storage volumes for all the views at cluster-2:

```
Vplexcli:/> export storage-view map
  /clusters/cluster-2/exports/storage-views/*
VPD83T3:6000144000000010a000e68dc5f76188 base01_vol
VPD83T3:6000144000000010a0014760d64cb21f dd_00_vol
VPD83T3:6000144000000010a0014760d64cb221 dd_01_vol
.
.
.
```

**See also**

- ◆ [export storage-view find-unmapped-volumes on page 258](#)
- ◆ [export storage-view find on page 257](#)
- ◆ [export storage-view summary on page 264](#)

## export storage-view removeinitiatorport

Removes the specified initiator-port(s) from the view.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **storage-view removeinitiatorport**.

In `/clusters/cluster/exports/storage-views` context, command is **removeinitiatorport**.

**Syntax** `export storage-view removeinitiatorport  
[-v|--view] context path  
[-i|--initiator-ports] context path,context path...`

**Arguments** **Required arguments**

`[-i|--initiator-ports] context path,context path...` - \* Comma-separated list of one or more initiator ports to remove.

**Optional arguments**

`[-v|--view] context path` - The storage view from which to remove the initiator port.

\* - argument is positional.

**Description** Use the `ll /clusters/cluster/exports/storage-views/storage-view` command to display the initiator ports in the specified storage view.

**Example** Remove an initiator port from `/clusters/cluster/exports/storage-views/storage-view` context:

```
Vplexcli:/clusters/cluster-1/exports/storage-views /LicoJ009>  
removeinitiatorport -i LicoJ009_hba1
```

**See also**

- ◆ [export storage-view addinitiatorport on page 250](#)
- ◆ [export storage-view removeport on page 261](#)

---

## export storage-view removeport

Removes the specified port(s) from a storage view.

**Contexts** All contexts.

In `/clusters/cluster/exports/storage-views/storage-view` context, command is **removeport..**

```
export storage-view removeport
  [-v|--view] context path
  [-p|--ports] context path,context path...
```

**Arguments** **Required arguments**

`[-p|--ports] context path,context path ...` - \* List of one or more ports to be removed from the view. Entries must be separated by commas.

**Optional arguments**

`[-v|--view] context path` - View from which to remove the specified port(s).

\* - argument is positional.

**Description** Use the `ll /clusters/cluster/exports/storage-views/storage-view` command to display the ports in the specified storage view

**Example** Remove a port from `/clusters/cluster/exports/storage-views/storage-view` context:

```
Vplexcli:/clusters/cluster-1/exports/storage-views/LicoJ009>
  removeport -p P000000003CA00147-A0-FC02
```

**See also**

- ◆ [export storage-view addport on page 251](#)
- ◆ [export storage-view destroy on page 256](#)

## export storage-view removevirtualvolume

Removes the specified virtual volume from the view.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **storage-view removevirtualvolume**.

In `/clusters/cluster/exports/storage-views` context, command is **removevirtualvolume**.

**Syntax**

```
export storage-view removevirtualvolume
  [-v|--view] context path
  [-f|--force]
  [-o|--virtual-volumes] volume, volume...
```

**Arguments** **Required arguments**

`[-o|--virtual-volumes] volume, volume ...` - \* List of one or more virtual volumes to be removed from the view. Entries must be separated by commas.

**Optional arguments**

`[-f|--force]` - Force the virtual-volumes to be removed from the view even if the specified LUNs are in use, the view is live, or some of the virtual-volumes do not exist in the view.

`[-v|--view] context path` - View from which to remove the specified virtual volume(s).

\* - argument is positional.

**Description** Use the `ll /clusters/cluster/exports/storage-views/storage-view` command to display the virtual volumes in the specified storage view

**Example** Delete a virtual volume from the specified storage view, even though the storage view is active:

```
Vplexcli:/clusters/cluster-1/exports/storage-views> removevirtualvolume --view E209_View --virtual-volume (1,test3211_r0_vol) --force
```

```
WARNING: The storage-view 'E209_View' is a live storage-view and is exporting storage through the following initiator ports: 'iE209_hba1_b', 'iE209_hba0'. Performing this operation may affect hosts' storage-view of storage. Proceeding anyway.
```

**See also**

- ◆ [export storage-view addvirtualvolume on page 252](#)
- ◆ [export storage-view destroy on page 256](#)

## export storage-view show-powerpath-interfaces

Displays the mapping between PowerPath® interfaces and the VPLEX system ports.

**Contexts** Clusters/*cluster* context and below.

**Syntax** `export storage-view show-powerpath-interfaces  
[-c|--cluster] context path`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-c|--cluster] context path` - The cluster at which to show the PowerPath interface mapping.

**Example**

```
Vplexcli:/clusters/cluster-2/exports/storage-views> export
storage-view show-powerpath-interfaces
PowerPath Interface Target Port
-----
SP A8                P000000003CA000E6-A0-FC03.0
SP A9                P000000003CB001CB-B0-FC01.0
SP A6                P000000003CA001CB-A0-FC01.0
SP A7                P000000003CB000E6-B0-FC01.0
SP A0                P000000003CA000E6-A0-FC00.0
SP A1                P000000003CA000E6-A0-FC01.0
SP A4                P000000003CA000E6-A0-FC02.0
SP A5                P000000003CB001CB-B0-FC00.0
```

**See also**

- ◆ [export storage-view checkconfig on page 254](#)
- ◆ [export storage-view find on page 257](#)
- ◆ [export storage-view map on page 259](#)
- ◆ [export storage-view summary on page 264](#)

## export storage-view summary

Lists each view and the number of virtual volumes and initiators that it contains.

**Contexts** All contexts.

In `/clusters/cluster/exports/storage-views` context, command is **summary**.

**Syntax** `export storage-view summary`  
`[-c|--clusters] cluster,cluster...`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-c|--cluster] cluster, cluster...` - List of clusters. Entries must be separated by commas. Display information only for storage views on the specified cluster(s).

**Description** At the root level, displays information for all clusters.

At the `/clusters/cluster` context and below, displays information only for views in the cluster in that context.

**Example** Display storage view summary for a specified cluster (no unhealthy views):

```
VPlexcli:/> export storage-view summary --clusters cluster-1
```

```
View health summary(cluster-1):
```

| view name | health-state | exported volumes | ports | registered initiators |
|-----------|--------------|------------------|-------|-----------------------|
| LicoJ009  | healthy      | 27               | 4     | 4                     |
| LicoJ013  | healthy      | 1                | 8     | 4                     |

```
Total 2 views, 0 unhealthy.
```

Display storage view summary for all clusters (1 unhealthy view):

```
VPlexcli:/> export storage-view summary
```

```
View health summary(cluster-1):
```

| view name  | health-state | exported volumes | ports | registered initiators |
|------------|--------------|------------------|-------|-----------------------|
| poly2_view | healthy      | 5                | 4     | 2                     |
| view1      | healthy      | 1                | 4     | 1                     |

```
Total 2 views, 0 unhealthy.
```

```
View health summary(cluster-2):
```

```
view name operational status port summary
```

| view name | health-state | port name                 | unhealthy volumes | export status |
|-----------|--------------|---------------------------|-------------------|---------------|
| esx1_view | error        | P000000003B2017D8-A0-FC00 | 1                 | ok            |
|           |              | P000000003B2017D8-A0-FC01 | 1                 | ok            |
|           |              | P000000003B3017D8-B0-FC00 | 1                 | ok            |
|           |              | P000000003B3017D8-B0-FC01 | 1                 | ok            |

| view name | health-state | exported volumes | ports | registered initiators |
|-----------|--------------|------------------|-------|-----------------------|
| esx1_view | unhealthy    | 10               | 4     | 2                     |

```
Total 1 views, 1 unhealthy.
```

- 
- See also**
- ◆ [export port summary on page 248](#)
  - ◆ [export storage-view checkconfig on page 254](#)
  - ◆ [export storage-view map on page 259](#)
  - ◆ [export storage-view show-powerpath-interfaces on page 263](#)
  - ◆ [storage-volume summary on page 469](#)

## export target-port renamewwns

Renames a target port's WWN pair.

**Contexts** All contexts.

In `/clusters/cluster/exports` context, command is **target-port renamewwns**.

**Syntax** `export target-port renamewwns  
[-p|--port] context path  
[-w|--wwns] wwns`

**Arguments** **Required arguments**

`[-w|--wwns] wwns` - A WWN pair separated by "|":

`portWWN|nodeWWN`

Each WWN is either '0x' followed by one or more hexadecimal digits or an abbreviation, in the following format:

`<string>:<number>[,<number>]`

For example,

`0xd1342a|0xd1342b`

`hyy1:194e,4|hyy1:194e`

`0xd1342a`

`hyy1:194e,4`

**Optional arguments**

`[-p|--port] context path` - Target port for which to rename the WWN pair.

**Description** Use the `ll` command in `/clusters/cluster/export/port` context to display portWWNs and nodeWWNs.



### CAUTION

**Disable the corresponding Fibre Channel port before executing this command.**

### Example

```
VPLexcli: /> export target-port renamewwns --wwns 0xd1342a|0xd1342b --port  
P0000000000000001-FK00
```

**See also** ♦ [export initiator-port discovery on page 243](#)

## extent create

Creates one or more storage-volume extents.

**Contexts** All contexts.

**Syntax**

```
extent create
  [-s|--size] size
  [-o|--block-offset] integer
  [-n|--num-extents] integer
  [-d|--storage-volumes] storage-volume,storage-volume...
```

**Arguments** **Required arguments**

**[-d|--storage-volumes] storage-volume,storage-volume ...** - \* Names of one or more claimed storage volumes to extent. Entries must be separated by commas.

**Optional arguments**

**[-s|--size] size** - The size of each extent, in bytes. If not specified, the largest available contiguous range of 4K byte blocks on the storage volume is used to create the specified number of extents.

**[-n|--num-extents] integer** - The number of extents to create per specified storage volume. Maximum of 128 extents per storage volume. If not specified, only one extent per storage-volume is created.

**[-o|--block-offset] integer** - The block-offset on the underlying storage volume on which the extent is created. If not specified, the block-offset is determined automatically.

\* - argument is positional.

**Description** An extent is a slice (range of 4K byte blocks) of a storage volume. An extent can use the entire capacity of the storage volume, or the storage volume can be carved into a maximum of 128 extents.

Extents are the building blocks for devices.

If the storage volume is larger than the desired virtual volume, create an extent the size of the desired virtual volume. Do not create smaller extents, and then use different RAID configurations to concatenate or stripe the extents.

If the storage volume is smaller than the desired virtual volume, create a single extent per storage volume, and then use devices to concatenate or stripe these extents into a larger device.

**Example** In the following example:

- ◆ The `ll -p **/storage-volumes` command displays a list of all storage volumes.
- ◆ The `cd` command changes the context to the storage-volume context on cluster-1.
- ◆ The `extent create` command creates an extent from two claimed 16 GB storage volumes.

```
Vplexcli:/> ll -p **/storage-volumes
```

| Name         | VPD83 ID                                  | Capacity | Use       | Vendor | IO    | Status      | Type        |
|--------------|-------------------------------------------|----------|-----------|--------|-------|-------------|-------------|
| CX4_Logging  | VPD83T3: 6006016021d02500e6d58bab2227df11 | 80G      | used      | DGC    | alive | normal      | normal      |
| CX4_Meta_M0  | VPD83T3: 6006016021d02500be83caff0427df11 | 90G      | claimed   | DGC    | alive | normal      | normal      |
| CX4_Meta_M1  | VPD83T3: 6006016021d02500bf83caff0427df11 | 90G      | claimed   | DGC    | alive | normal      | normal      |
| CX4_lun0     | VPD83T3: 6006016021d0250026b925ff60b5de11 | 10G      | used      | DGC    | alive | normal      | normal      |
| Log1723_154  | VPD83T3: 60000970000192601723533030313534 | 78G      | used      | EMC    | alive | normal      | normal      |
| Log1852_154  | VPD83T3: 60000970000192601852533030313534 | 78G      | used      | EMC    | alive | normal      | normal      |
| Metal723_150 | VPD83T3: 60000970000192601723533030313530 | 78G      | meta-data | EMC    | alive | traditional | traditional |
| Metal852_150 | VPD83T3: 60000970000192601852533030313530 | 78G      | meta-data | EMC    | alive | traditional | traditional |
| Symm1723_1C0 | VPD83T3: 60000970000192601723533030314330 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1C4 | VPD83T3: 60000970000192601723533030314334 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1C8 | VPD83T3: 60000970000192601723533030314338 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1CC | VPD83T3: 60000970000192601723533030314343 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1D0 | VPD83T3: 60000970000192601723533030314430 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1D4 | VPD83T3: 60000970000192601723533030314434 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1D8 | VPD83T3: 60000970000192601723533030314438 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1DC | VPD83T3: 60000970000192601723533030314443 | 16G      | claimed   | EMC    | alive | normal      | normal      |
| Symm1723_1E0 | VPD83T3: 60000970000192601723533030314530 | 16G      | claimed   | EMC    | alive | normal      | normal      |

```
VPlexcli: />cd /clusters/cluster-1/storage-elements/storage-volumes
```

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> extent create  
Symm1723_1DC, Symm1723_1E0
```

- See also**
- ◆ [extent create on page 267](#)
  - ◆ [extent destroy on page 269](#)

---

## extent destroy

Destroys one or more storage-volume extents.

**Contexts** All contexts.

**Syntax** `extent destroy`  
`[-f|--force]`  
`[-s|--extents context path,context path...]`

**Arguments** **Required arguments**  
`[-s|--extents] context path,context path ...` - \* List of one or more extent(s) to destroy. Entries must be separated by commas.

**Optional arguments**  
`[-f|--force]` - Forces the destruction of the given extents, bypassing all guards and confirmations.

\* - argument is positional.

**Description** Destroys the specified extents.

**Example** Destroy an extent:

```
Vplexcli:/clusters/cluster-1/storage-elements/extent> extent destroy --force  
extent_Symm1254_7BA_1  
Destroyed 1 out of 1 targeted extents.
```

**See also** ♦ [extent create on page 267](#)

## extent summary

Displays a list of a cluster's unhealthy extents.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/extents` context, command is **summary**.

**Syntax** `extent summary`  
`[-c|--clusters] cluster,cluster...`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-c|--clusters] cluster,cluster...` - List of clusters to summarize, separated by commas. May contain glob characters.

**Description** Displays a cluster's unhealthy extents (if any exist), the total number of extents by use, and calculates the total extent capacity for this cluster.

An unhealthy extent has a non-nominal health state, operational status or I/O status.

If the `--clusters` argument is not specified and the command is executed at or below a specific cluster's context, information is summarized for only that cluster. Otherwise, the extents of all clusters are summarized.

**Example** Output sample (partial):

```
Vplexcli:/> extent summary
SUMMARY (cluster-1)
Extent Name          IO Status  Operational Status  Health State
-----
extent_Symm1723_200_1  alive     error               critical-failure
extent_Symm1723_204_1  alive     error               critical-failure
extent_Symm1723_208_1  alive     error               critical-failure
.
.
.
Extent Summary
-----

Health      extents      150
            out-of-date   0
            unhealthy  40

Use         used         138
            claimed    5
            unclaimed  0
            unusable   5
            logging   2

Capacity   total       2.43T

SUMMARY (cluster-2)
```

**Table 15 extent summary field descriptions**

| Field                                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Health summary (displayed only for unhealthy extents)</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Name                                                         | Name of extent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| I/O Status                                                   | <b>alive</b> - I/O is proceeding normally on the extent.<br><b>dead</b> - The underlying storage volume is marked as hardware-dead.<br><b>unreachable</b> - The underlying storage volume is unreachable.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Operational Status                                           | <b>degraded</b> - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).<br><b>ok</b> - The extent is functioning normally.<br><b>starting</b> - The extent is not yet ready.<br><b>unknown</b> - VPLEX cannot determine the extent's Operational state, or the state is invalid.                                                                                                                                                                                                                                                                                                                           |
| Health State                                                 | <b>degraded</b> - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).<br><b>ok</b> - The extent is functioning normally.<br><b>non-recoverable-error</b> - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device), and/or the Health state cannot be determined.<br><b>unknown</b> - VPLEX cannot determine the extent's Operational state, or the state is invalid.                                                                                                                                                                            |
| <b>Extent Summary</b>                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Health                                                       | <b>extents</b> - Total number of extents on the cluster, the number.<br><b>out-of-date</b> - Of the total number of extents on the cluster, the number that are out-of-date compared to their mirror.<br><b>unhealthy</b> - Of the total number of extents on the cluster, the number with operational status or health state that is not "ok".                                                                                                                                                                                                                                                                                                                       |
| Use                                                          | <b>used</b> - Of the total number of extents on the cluster, the number in use.<br><b>claimed</b> - Of the total number of extents on the cluster, the number that are claimed<br><b>unclaimed</b> - Of the total number of extents on the cluster, the number that are unclaimed.<br><b>unusable</b> - Indicates that the underlying storage-volume of the extent is dead or unreachable. Use the "storage-volume summary" command to check the storage-volume. Use the "validate-system-configuration" command to check reachability from the directors.<br><b>logging</b> - Of the total number of extents on the cluster, the number that are in use for logging. |
| Capacity                                                     | Total capacity on the cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- See also**
- ◆ [ds summary on page 234](#)
  - ◆ [export port summary on page 248](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [local-device summary on page 283](#)
  - ◆ [storage-volume summary on page 469](#)
  - ◆ [virtual-volume summary on page 510](#)

## getsysinfo

Returns information about the current system.

**Contexts** All contexts.

**Syntax** getsysinfo  
--output *path name*  
--linux

**Arguments** **Required arguments**  
None.

**Optional arguments**

**--output *path name*** - Location and name of the output file.

Default: /var/log/VPLEX/cli/YY-sysinfo.txt

**--linux** - Use this if the management server is running on a Linux system. Disables the scsi tests since Linux systems lack a scsi command.

**Description** The information is written in TCL format.

**Example** Display information and send the output to a file:

```
VPLEXcli: /> getsysinfo --output /var/log/VPLEX/cli/TestGetSysInfo
Running from localhost:/var/log/VPLEX/cli
Local time 2010/08/04 14:22:43
Flag includeDebug = 0
Flag isLinux = 0
Treating this tower like version D4
clustercount = 2
26 ports - unknown system type
System does NOT have comtcp enabled

## ===== NCB: Cluster_1_Dir1A =====
## ===== NCB: Cluster_1_Dir1B =====
.
.
.
Raw output saved to /var/log/VPLEX/cli/TestGetSysInfo.2010-08-04-14.25.raw
```

**Table 16** getsysinfo field descriptions

| Field                               | Description                                                                                                             |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Flag includeDebug                   |                                                                                                                         |
| Flag isLinux = 0                    | Denotes a linux simulator build. Ignore this line.                                                                      |
| Treating this tower like version D4 | Denotes the system is Release 4.0 or later. Ignore this line.                                                           |
| nn ports - unknown system type      | The getsysinfo script looked for hardware prior to Release 4.0 and did not find it.                                     |
| System does NOT have comtcp enabled | Communication protocol used on Ethernet ports for connections to other clusters prior to Release 4.0. Ignore this line. |

- See also**
- ◆ [cluster summary on page 93](#)
  - ◆ [director firmware show-banks on page 197](#)
  - ◆ [manifest version on page 302](#)
  - ◆ [version on page 501](#)

## health-check

Displays overall hardware/software health.

**Contexts** All contexts.

**Syntax** health-check  
[-m|--highlevel]  
[-f|--full]

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-m|--highlevel]** - Checks for major subcomponents with error conditions, warnings or ok. Used for instantaneous, high level view of the health of the VPLEX.

Default behavior if no other argument is specified.

**[-f|--full]** - Runs full scan.

**Description** High level view of the health of the VPLEX.

Consolidates information from the following commands:

- ◆ version
- ◆ cluster status
- ◆ cluster summary
- ◆ connectivity validate-be
- ◆ connectivity validate-wan-com
- ◆ ds summary
- ◆ export storage-view summary
- ◆ virtual-volume summary
- ◆ storage-volume summary
- ◆ ll /clusters/\*\*/system-volumes/

**Example** Run a high-level (default) health check on a VPLEX Metro:

```
Vplexcli: /> health-check
Product Version: 5.1.0.00.00.10

Clusters:
-----
Cluster   Cluster  Oper   Health   Connected  Expelled
Name      ID       State  State
-----
cluster-1 1        ok     degraded True        False
cluster-2 2        ok     ok       True        False

cluster-1 Transition/Health Indications:
Device initializing
20 unhealthy Devices or storage-volumes

Meta Data:
-----
Cluster   Volume   Volume   Oper   Health   Active
Name      Name     Type     State  State
-----
cluster-1 Advil_1  meta-volume ok     ok       True
cluster-1 logging_c1_log_vol logging-volume ok     ok       -
cluster-1 Advil_1_backup_2012Mar07_043012 meta-volume ok     ok       False
cluster-1 Advil_1_backup_2012Mar08_043011 meta-volume ok     ok       False
cluster-2 logging_c2_log_vol logging-volume ok     ok       -
```

```

cluster-2 Advil-2_backup_2012Mar08_043020 meta-volume ok ok False
cluster-2 Advil-2_backup_2012Mar07_043017 meta-volume ok ok False
cluster-2 Advil-2 meta-volume ok ok True

```

Front End:

```

-----
Cluster Name Total Storage Views Unhealthy Storage Views Total Registered Initiators Total Ports Total Exported Volumes Total ITLs
-----
cluster-1 4 2 12 8 135 672
cluster-2 0 0 0 0 0 0

```

Storage:

```

-----
Cluster Name Total Storage Volumes Unhealthy Storage Volumes Total Virtual Volumes Unhealthy Virtual Volumes Total Dist Devs Unhealthy Dist Devs No Dual Paths Not visible from All Dirs
-----
cluster-1 2375 10 229 10 12 0 0 0
cluster-2 2365 0 205 0 12 0 0 0

```

Consistency Groups:

```

-----
Cluster Name Total Synchronous Groups Unhealthy Synchronous Groups Total Asynchronous Groups Unhealthy Asynchronous Groups
-----
cluster-1 9 0 0 0
cluster-2 5 0 0 0

```

FC WAN Connectivity:

```

-----
Port Group Connectivity
-----
port-group-1 ok
port-group-0 ok

```

Cluster Witness:

Cluster Witness is not configured

RecoverPoint:

```

-----
Cluster Name Total RP Clusters Unhealthy RP Clusters Total Replicated Virtual Volumes Unhealthy Replicated Virtual Volumes Total RP-enabled Consistency Groups Mis-aligned RP-enabled Consistency Groups Total Registered RP Initiators/Storage Views Unhealthy RP Storage Views
-----
cluster-1 1 1 1 0 8 1 8/1 0
cluster-2 - - - - - - -/- -

```

\*\*This command is only able to check the health of the local cluster(cluster-1)'s RecoverPoint configuration, therefore if this system is a VPLEX Metro or VPLEX Geo repeat this command on the remote cluster to get the health of the remote cluster's RecoverPoint configuration.

**Example** Run a full-scan health-check on a VPLEX Metro:

```

Vplexcli:/> health-check --full
Configuration (CONF):
Checking VplexCli connectivity to directors..... OK
Checking Directors Commission..... OK
Checking Directors Communication Status..... OK
Checking Directors Operation Status..... OK
Checking ports status..... Error
Checking Call Home..... Warning
Checking Connectivity..... OK
Back End (BE):
Checking Unreachable Storage Volumes..... Error
Checking Unhealthy Virtual Volumes..... Error
Error devices in SymmProd_Vol_0000_vol

/clusters/cluster-1/devices/SymmProd_Vol_0000:

Attributes:
Name Value
-----

```

```

application-consistent  false
auto-resume              -
block-count              6576000
block-offset             0
block-size               4K
capacity                 25.1G
geometry                 raid-0
health-indications      []
health-state             critical-failure
locality                  local
operational-status      error
rebuild-allowed         -
rebuild-eta              -
rebuild-progress        -
rebuild-status          -
rebuild-type            -
rule-set-name           -
service-status          -
stripe-depth            -
system-id                SymmProd_Vol_0000
transfer-size           -
virtual-volume           SymmProd_Vol_0000_vol
visibility               local

```

Contexts:

| Name       | Description                                                               |
|------------|---------------------------------------------------------------------------|
| components | The list of components that support this device or system virtual volume. |

```

/clusters/cluster-1/devices/SymmProd_Vol_0000/components/SymmProd_Vol_0000_extent_0:
.
.
.
/clusters/cluster-1/devices/SymmRep_Vol_0004:
.
.
.
/clusters/cluster-1/devices/SymmRep_Vol_0004/components/SymmRep_Vol_0004_extent_0:
.
.
.
Back end array status..... Error
Validating paths to back end arrays..... Error
Front End (FE):
Checking Storage Views..... Error
Checking Front End Path..... Warning
Cache:
Checking for sub-pages writes(25% or above of total writes).... Error
Checking Stuck IO..... OK
Director Health Status:
Checking Director Hardware..... Error
Checking SSD Hardware..... OK
Checking Director RPMs..... OK

```

Output to /var/log/VPlex/cli/health\_check\_full\_scan.log

- See also**
- ◆ [cluster status on page 89](#)
  - ◆ [validate-system-configuration on page 492](#)

## help

Displays help on one or more commands.

**Contexts** All contexts.

**Syntax** help  
[-i|--interactive]  
[-G|--no-global]  
[-n|--no-internal]

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-i|--interactive]** - Invoke interactive help.

**[-G|--no-global]** - Suppresses the list of global commands for contexts other than root context.

**[-n|--internal]** - Include commands that are normally used for low-level debugging and development.

Type **quit** to exit interactive help.

**Description** If an argument is marked as required, it is always required. Additional arguments may be required depending on the context in which the command is executed.

**Example** Display only commands specific to the current context:

```
VPlexcli:/clusters/cluster-1> help -G  
Commands inherited from parent contexts:  
add cacheflush configdump expel forget shutdown summary unexpel
```

```
Commands specific to this context and below:  
status verify
```

Display help for a specified command:

```
VPlexcli:/clusters/cluster-1> status --help  
synopsis: status [<options>]
```

Displays a cluster's operational-status and health-state.

options (\* = required):

```
-h, --help  
    Displays the usage for this command.  
--verbose  
    Provide more output during command execution. This may not have any effect for some  
commands.  
-c, --clusters= <clusters>  
    clusters whose operational-status to display.
```

Along with the operational-status, an indication of why it could be non-nominal and a progress indicator are displayed.

Health-state has a similar indicator.

Invoke interactive help:

```
VPlexcli:/clusters/cluster-1> help -i  
  
Welcome to Python 2.2! This is the online help utility.  
. . .  
help> topics
```

Here is a list of available topics. Enter any topic name to get more help.

|            |                 |               |            |
|------------|-----------------|---------------|------------|
| ASSERTION  | DYNAMICFEATURES | NONE          | TRACEBACKS |
| ASSIGNMENT | ELLIPSIS        | NUMBERMETHODS | TRUTHVALUE |

.  
.  
.

help> **EXPRESSIONS**

-----

#### 5.14 Summary

The following table summarizes the operator precedences in Python, from lowest precedence (least binding) to highest precedence (most binding).

.  
.  
.

## history

Displays or clears the command history list.

**Contexts** All contexts.

**Syntax** `history`  
`[-c|--clear]`  
`[-n|--number] number`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-c|--clear]` - Clears the history list.

`[-n|--number] number` - Displays only the last *n* commands in the history list.

**Example** Display the last 8 commands executed in this CLI session:

```
VPlxcli:/> history 8
492 ll
493 cd d
494 cd device-migrations/
495 ll
496 cd
497 ds summary
498 export storage-view checkconfig
499 history 8
```

## local-device create

Creates a new local-device.

**Contexts** All contexts.

**Syntax**

```
local-device create
[-d|--stripe-depth] depth
[-n|name] name
[-g|--geometry] {raid-0|raid-1|raid-c}
[-e|extents] context path,context path...
[-s|--source-leg] context path
--force
```

**Arguments** **Required arguments**

**[-n|--name] name** - \* Name for the new device. Must be unique across all clusters. Devices on different clusters that have the same name cannot be combined into a distributed device.

**Note:** If this device will have another device attached (using the “[device attach-mirror](#)” command to create a RAID-1), the name of the resulting RAID-1 is the name given here plus a timestamp. Names in VPLEX are limited to 63 characters. The timestamp consumes 16 characters. Thus, if this device is intended as the parent device of a RAID-1, the device name must not exceed 47 characters.

**[-g|--geometry] {raid-0|raid-1|raid-c}** - \* Geometry for the new device. Valid values are “raid-0”, “raid-1”, or “raid-c”.



### CAUTION

**Use this command to create a RAID 1 device only if:**

- None of the legs contains data that must be preserved,
- The resulting device will be initialized using tools on the host, or
- The resulting device will be added as a mirror to another device.

**[-e|--extents] context path,context path...** - \* List of one or more claimed extents to be added to the device. Can also be other local devices (to create a device of devices).

### Optional arguments

**[-d|--stripe-depth] depth** - Required if **--geometry** is raid-0. Stripe depth must be:

- ◆ Greater than zero
- ◆ No greater than the number of blocks of the smallest element of the RAID 0 device being created
- ◆ A multiple of the block size: 4 K bytes

A depth of 32 means 128 K (32 x 4 K) is written to the first disk then the next 128 K is written to the next disk.

Concatenated RAID devices are not striped.

**[-s|--source-leg] context path** - When geometry argument is raid-1, picks one of the extents specified by the **--extents** argument to be used as the source data image for the new device. The command copies data from the **--source-leg** to the other legs of the new device.

**[-f|--force]** - Create a raid-1 device even if no **--source-leg** is specified.

\* - argument is positional.

## Description

A device is configured from one or more extents in a RAID 1, RAID 0, or concatenated (RAID C) configuration.

The block sizes of the supporting extents must be the same (4 K bytes) and will determine the local-device block size.

When creating a device with RAID 1 geometry, this command prints a warning and asks for confirmation.



### WARNING

*If the `--source-leg` argument is not specified, this command does not initialize or synchronize the legs of a RAID 1 device. Because of this, a RAID 1 device created by this command does not guarantee that consecutive reads of the same block will return the same data if the block has never been written.*

To create a RAID 1 device when one leg of the device contains data that must be preserved, use the `--source-leg` argument or the `“device attach-mirror”` command to add a mirror to the leg.

By default, automatic device rebuilds are enabled on all devices. For configurations with limited bandwidth between clusters, it may be useful to disable automatic rebuilds.

Use the `“set”` command to enable/disable automatic rebuilds on the distributed device. The rebuild setting is immediately applied to the device.

- ◆ Set `rebuild-allowed` to **true** to start or resume a rebuild if the mirror legs are out of sync.
- ◆ Set `rebuild-allowed` set to **false** to stop any rebuild in progress.

When automatic rebuild is re-enabled on a device where it has been disabled, the rebuild starts again from the place where it stopped.

## Example

In the following example:

- ◆ The `ll` command displays the available (claimed) extents
- ◆ The `local-device create` command is used to create a 16 GB RAID 1 device named `TestDevCluster1` on cluster 1
- ◆ The `cd` command returns to the root context
- ◆ The `ll -p **/devices` command displays the new device

```
Vplexcli:/clusters/cluster-1/storage-elements/extents> ll
```

| Name                  | StorageVolume | Capacity | Use     |
|-----------------------|---------------|----------|---------|
| .....                 | .....         | .....    | .....   |
| .                     |               |          |         |
| .                     |               |          |         |
| .                     |               |          |         |
| extent_Symm1852_AAC_1 | Symm1852_AAC  | 16G      | claimed |
| extent_Symm1852_AB0_1 | Symm1852_AB0  | 16G      | claimed |
| extent_Symm1852_AB4_1 | Symm1852_AB4  | 16G      | claimed |
| extent_Symm1852_AB8_1 | Symm1852_AB8  | 16G      | claimed |

```
Vplexcli:/clusters/cluster-1/storage-elements/extents> local-device create --name
TestDevCluster1 --geometry raid-1 --extents
/clusters/cluster-1/storage-elements/extents/extent_Symm1852_AAC_1,/clusters/cluster-1/storag
e- elements/extents/extent_Symm1852_AB0_1
```

```
Vplexcli:/clusters/cluster-2/storage-elements/extents> cd
Vplexcli:/> ll -p **/devices
```

```

/clusters/cluster-1/devices:
Name           Operational Health  Block   Block   Capacity  Geometry  Visibility  Transfer  Virtual
----- Status   State  Count   Size   -----  -----  -----  Size     Volume
-----
TestDevCluster1 ok      ok      4195200 4K      16G      raid-1     local      2M        -
base0          ok      ok      262144  4K      1G       raid-0     local      -         base0_vol
base1          ok      ok      262144  4K      1G       raid-0     local      -         base1_vol

```

In the following example, the **local-device create** command creates a RAID-1 device from 2 extents; *extent\_lun\_1\_1* and *extent\_lun\_2\_1* in which:

- ◆ *extent\_lun\_2\_1* is the same size or larger than *extent\_lun\_1\_1*
- ◆ *extent\_lun\_1\_1* is the source leg of the new device
- ◆ *extent\_lun\_2\_1* is the mirror leg

```

Vplexcli: /> local-device create --geometry raid-1 --extents extent_lun_1_1, extent_lun_2_1
--name dev_lun_1 --source-leg extent_lun_1_1

```

```

Vplexcli: /> ls -al /clusters/cluster-1/devices/

```

```

/clusters/cluster-1/devices:
Name           Operational Health  Block   Block   Capacity  Geometry  Visibility  Transfer  Virtual
----- Status   State  Count   Size   -----  -----  -----  Size     Volume
-----
dev_lun_1      ok      ok      20709376 4K      5G       raid-1     local      -         -

```

- See also**
- ◆ [device attach-mirror on page 181](#)
  - ◆ [local-device destroy on page 282](#)
  - ◆ [local-device summary on page 283](#)

## local-device destroy

Destroys existing local-devices.

**Contexts** All contexts.

**Syntax** `local-device destroy`  
`[-f|--force]`  
`[-d|--devices] context path,context path...`

**Arguments** **Required arguments**  
`[-d|--devices] context path,context path...` - \* List of one or more device(s) to destroy.

**Optional arguments**  
`[-f|--force]` - Force the destruction of the devices without asking for confirmation.  
\* - argument is positional.

**Description** The device must not be hosting storage or have a parent device.

**Example** `Vplexcli:/clusters/cluster-1> local-device destroy -d was_1_leg_r1`

WARNING: The following items will be destroyed:

```
Context
-----
/clusters/cluster-1/devices/was_1_leg_r1
Do you wish to proceed? (Yes/No)
```

**See also** ♦ [local-device create on page 279](#)

## local-device summary

Displays unhealthy local devices and a summary of all local devices.

**Contexts** All contexts.

In `/clusters/cluster/devices` context, command is **summary**.

**Syntax** `local-device summary`  
`[-c|--clusters] cluster,cluster...`

**Arguments** **Required arguments**  
 None.

**Optional arguments**

`[-c|--clusters] cluster,cluster...` - Display information only for the specified cluster(s).

**Description** Displays unhealthy local devices and a summary of all local devices. Unhealthy devices have non-nominal health state, operational status, or service-status.

If the `--clusters` argument is not specified and the command is executed at or below a `/clusters/cluster` context, information for only that cluster is displayed.

**Example** Display local devices for a specified cluster:

```
Vplexcli:/> local-device summary --clusters cluster-1
device name      health state      operational status  service status
-----
dev_sym1723_1FC  critical-failure  error              suspended

Device Summary  (no tier)
-----
Health          devices           5
                unhealthy        1

Visibility      local             5

Capacity        devices w/ space  0
                free capacity     0B
                total capacity   12G
```

**Table 17** local device summary field descriptions

| Field             | Description                                                                                                                                                                                                                                                                                                                                                 |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Health</b>     |                                                                                                                                                                                                                                                                                                                                                             |
| devices           | Number of devices in the cluster.                                                                                                                                                                                                                                                                                                                           |
| unhealthy         | Of the total number of devices in the cluster, the number whose health state is not "ok".                                                                                                                                                                                                                                                                   |
| <b>Visibility</b> | Of the total number of devices in the cluster, the number with global or local visibility.<br><b>global</b> - The remote cluster can access the virtual volume. A virtual volume on a top-level device that has global visibility can be exported in storage views on any cluster.<br><b>local</b> (default) - Device is visible only to the local cluster. |
| <b>Capacity</b>   |                                                                                                                                                                                                                                                                                                                                                             |
| devices w/ space  | Of the total number of devices in the cluster, the number with available space.                                                                                                                                                                                                                                                                             |
| free capacity     | Total free capacity on the cluster.                                                                                                                                                                                                                                                                                                                         |
| total capacity    | Total capacity of the cluster.                                                                                                                                                                                                                                                                                                                              |

- See also**
- ◆ [ds summary on page 234](#)
  - ◆ [export port summary on page 248](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [extent summary on page 270](#)
  - ◆ [storage-volume summary on page 469](#)
  - ◆ [virtual-volume summary on page 510](#)

## log filter create

Adds a new firmware log filter.

**Contexts** All contexts.

**Syntax**

```
log filter create
  [-s|--source] id
  [-t|--threshold] [<|>|=]0 - 7
  [-c|--component] name
  [-e|--event-num] id
  [-m|--message] text
  [-n|--no-consume]
```

**Arguments** **Required arguments**  
None.

### Optional arguments

**[-s|--source] *id*** - ID of the source log to be filtered. Use the “[log source list](#)” command to display the list of source logs and their IDs.

**[-t|--threshold] [*<|>|=*]0 - 7** - Severity of the events to write to the new log. Messages are categorized into 8 severities (0 - 7), with 0 being the most severe:

7 - **debug** (debug-level messages)

6 - **info** (informational messages)

5 - **notice** (normal but significant messages)

4 - **warning** (warning messages)

3 - **err** (error messages)

2 - **crit** (critical messages)

1 - **alert** (messages that must be handled immediately)

0 - **emerg** (messages notifying the system as unusable)

Default modifier is ‘>’.

**[-c|--component] *name*** - Component name to filter. Takes a regular expression as an argument. Plain strings are searched for in the component name.

**[-e|--event-num] *id*** - Used in conjunction with a specified component. An event ID to filter.

**[-m|--message] *text*** - An expression to look for in the event message. Takes a regular expression as an argument. Plain strings are searched for in the message text.

**[-n|--no-consume]** - Do not halt event processing after an event matches a filter.

**Description** Log filters define criteria for the destination of specific log data. A filter is placed in an ordered list, and filters see a received event in the order they sit in the list (shown by the **log filter list** command).

By default, filters consume received events so that a matching filter stops the processing of the event. Use the **--no-consume** argument to create a filter that allows processing of matching events to continue.

**Example** Filter out (hide) all messages with the string *test* in them:

```
Vplexcli: /> log filter create -m "test"
Filter added.
```

Filter all messages into the events log generated by the logserver component with the string Test:

```
Vplexcli:/> log filter create --source 1 --component logserver --message Test  
Filter added.
```

```
Vplexcli:/> log filter list  
1. [Source='/var/log/Vplex/cli/events.log', Component='logserver', Message matches 'Test']  
Destination='null' Consume='true'  
2. Component='logserver' Destination='null' Consume='true'  
3. [Threshold='>0'] Destination='null' Consume='true'
```

- See also**
- ◆ [log filter destroy on page 287](#)
  - ◆ [log filter list on page 288](#)

---

## log filter destroy

Removes a firmware log filter.

**Contexts** All contexts.

**Syntax** `log filter destroy  
[-f|--filter] filter`

**Arguments** **Required arguments**  
`[-f|--filter] filter` - ID of filter to be deleted.

**Description** The filter is removed from the filter stack.  
Use the “[log filter list](#)” command to display the filters configured on the VPLEX, and their associated IDs.

### Example

```
VPlexcli:/> log filter list
```

```
1. [Source='/var/log/VPlex/cli/events.log', Component='logserver', Message matches 'Test']  
   Destination='null' Consume='true'  
2. Component='logserver' Destination='null' Consume='true'  
3. [Threshold='>0'] Destination='null' Consume='true'
```

```
VPlexcli:/> log filter destroy 1  
Filter removed.
```

**See also**

- ◆ [log filter create on page 285](#)
- ◆ [log filter list on page 288](#)

## log filter list

Lists firmware log filters, in the order that they see events.

**Contexts** All contexts.

**Syntax** log filter list

**Description** The number printed beside each filter serves as both an identifier for the **log filter destroy** command as well as the order in which each respective filter will see an event.

### Example

```
VPLEXcli:/> log filter list
```

1. [Message matches 'Family and Fru Id Mismatch Retrieved'] Destination='null' Consume='true'
2. [Component='logserver'] Destination='null' Consume='true'
3. [Threshold='>=4'] Destination='null' Consume='true'

**See also** ♦ [log filter create on page 285](#)

♦ [log filter destroy on page 287](#)

## log source create

Adds a firmware log source.

**Contexts** All contexts.

**Syntax**

```
log source create
[-s|--source] host:port
[-p|--password password
[-f|--failover-source] host:port
```

**Arguments** **Required arguments**  
[-s|--source] *host:port* - \* IP address and port of the log source to be added. IP addresses of the VPLEX hardware components are listed in the *VPLEX Installation and Setup Guide*.

[-p|--password] *password* - The password to use for authenticating to the source.

**Optional arguments**

[-f|--failover-source] *host:port* - IP address and port of the failover source to be added.

\* argument is positional.

### Description



#### **CAUTION**

**For use by EMC personnel only.**

Creates a source for writing entries to the firmware log.

**Example** VPlxcli:/> **log source create --source 128.221.252.69:5988**

Enter the source connection password:

```
VPlxcli:/> log source list
1. /var/log/VPlx/cli/events.log
.
.
6. [128.221.252.69:5988]/cpu0/log
7. [128.221.252.69:5988]/xmmg/log
```

**See also**

- ◆ [log source destroy on page 290](#)
- ◆ [log source list on page 291](#)

## log source destroy

Destroys the specified log source.

**Contexts** All contexts.

**Syntax** `log source destroy  
[-s|--source] host:port`

**Arguments** **Required arguments**  
`[-s|--source] host:port` - IP address and port of the log source to destroy. IP addresses of the VPLEX hardware components are listed in the *VPLEX Installation and Setup Guide*.

### Description



#### **CAUTION**

**For use by EMC personnel only.**

### Example

```
VPlxcli:/> log source list
1. /var/log/VPlx/cli/events.log
2. 128.221.252.67:5988, [128.221.253.67:5988]/cpu0/log
3. 128.221.252.67:5988, [128.221.253.67:5988]/xmmg/log
4. 128.221.253.68:5988, [128.221.252.68:5988]/cpu0/log
5. 128.221.253.68:5988, [128.221.252.68:5988]/xmmg/log
6. [128.221.252.69:5988]/cpu0/log
7. [128.221.252.69:5988]/xmmg/log
8. [128.221.252.70:5988], 128.221.253.70:5988/cpu0/log
9. [128.221.252.70:5988], 128.221.253.70:5988/xmmg/log
```

```
VPlxcli:/> log source destroy --source 128.221.252.69:5988
```

**See also**

- ◆ [log source create on page 289](#)
- ◆ [log source list on page 291](#)

---

## log source list

Lists the various log paths from which log events are processed.

**Contexts** All contexts.

**Syntax** `log source list`

**Description** Lists the log paths from which log events are processed and their reference IDs.  
Used to create log filters.

**Example**

```
Vplexcli:/> log source list
1. /var/log/Vplex/cli/events.log
2. 128.221.252.35:5988,[128.221.253.35:5988]/xmmg/log
3. 128.221.252.36:5988,[128.221.253.36:5988]/cpu0/log
4. [128.221.252.35:5988],[128.221.253.35:5988]/cpu0/log
5. [128.221.252.36:5988],[128.221.253.36:5988]/xmmg/log
```

**See also**

- ◆ [log filter create on page 285](#)
- ◆ [log source create on page 289](#)

## logging-volume add-mirror

Adds a logging volume mirror.

**Contexts** All contexts.

**Syntax** `logging-volume add-mirror`  
`[-v|--logging-volume] logging volume`  
`[-m|--mirror] {name|context path}`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-v|--logging-volume] logging-volume` - Logging volume to which to add the mirror.

`[-m|--mirror] {name|context path}` - The name or context path of the device or storage-volume extent to add as a mirror. Must be top-level device or a storage-volume extent.

**Example** In the following example:

- ◆ The `ll` command in system-volumes context displays the names of existing logging volumes.
- ◆ The `logging-volume add-mirror` command adds a mirror to one of the logging volumes.

```
Vplexcli:/clusters/cluster-1/system-volumes> ll
Name      Volume Type  Operational Health  Active  Ready  Geometry  Block  Block  Capacity  Slots
-----  -
log1_vol  logging-volume ok      ok      -      -      raid-1    20448000 4K     78G      -
metal     meta-volume  ok      ok      true   true   raid-1    23592704 4K     90G     32000
```

```
Vplexcli:/clusters/cluster-1/system-volumes> logging-volume add-mirror --logging-volume
log1_vol --mirror /clusters/cluster-2/storage-elements/extends/extent_Symm2194_20F0_1
```

- See also**
- ◆ [logging-volume create on page 293](#)
  - ◆ [logging-volume destroy on page 295](#)

## logging-volume create

Creates a new logging volume in a cluster.

**Contexts** All contexts.

**Syntax**

```
logging-volume create
[-n|--name] name
[-g|--geometry {raid-0 |raid-1}
[-e|--extents] context path,context path...
[-d|--stripe-depth] depth
```

### Arguments **Required arguments**

**[-n|--name] name** - \* Name for the new logging volume.

**[-g|--geometry] {raid-0 |raid-1}** - \* Geometry for the new volume.

**[-e|--extents] context path,context path...-** \* List of one or more storage-volume extents to use to create the logging volume. Must not be empty, and must contain storage-volume extents that are all at the specified cluster. Entries must be separated by commas.

### **Optional arguments**

**[-d|--stripe-depth] depth** - Required if --geometry is raid-0. Stripe depth must be:

- ◆ Greater than zero, but not greater than the number of blocks of the smallest element of the RAID 0 device being created
- ◆ A multiple of 4 K bytes

A depth of 32 means 128 K (32 x 4 K) is written to the first disk, then the next 128 K is written to the next disk.

Best practice regarding stripe depth is to follow the best practice of the underlying array.

Concatenated RAID devices are not striped.

\* - argument is positional.

**Description** Creates a logging volume. The new logging volume is immediately available for use with distributed-devices.

A logging volume is required on each cluster in VPLEX Metro and Geo configurations. Each logging volume must be large enough to contain one bit for every page of distributed storage space (approximately 10 GB of logging volume space for every 160 TB of distributed devices).

Logging volumes experience a large amount of I/O during and after link outages. Best practice is to stripe each logging volume across many disks for speed, and to have a mirror on another fast disk.

To create a logging volume, first claim the storage volumes that will be used, and create extents from those volumes.

- ◆ Use the **ll /clusters/cluster/storage-elements/storage-volumes** command to display the available storage volumes on the cluster.
- ◆ Use the **"storage-volume claim" -n storage-volume\_name** command to claim one or more storage volumes.
- ◆ Use the **"extent create" -d storage-volume\_name, storage-volume\_name** command to create an extent to use for the logging volume.

Repeat this step for each extent to be used for the logging volume.

**Example** `Vplexcli:/clusters/cluster-2/system-volumes> logging-volume create -n cluster_2_log_vol -g raid-1 -e extent_1 , extent_2`

- See also**
- ◆ [extent create on page 267](#)
  - ◆ [logging-volume add-mirror on page 292](#)
  - ◆ [logging-volume destroy on page 295](#)
  - ◆ [storage-volume claim on page 458](#)

---

## logging-volume destroy

Destroys an existing logging volume.

**Contexts** All contexts.

**Syntax** `logging-volume destroy logging-volume`

**Arguments** **Required arguments**  
`[-v | --logging-volume] logging-volume - *` Name of logging volume to destroy.  
\* - argument is positional.

**Description** The volume to be destroyed must not be currently used to store block write logs for a distributed-device.

**Example** `Vplexcli:/clusters/cluster-1/system-volumes> logging-volume destroy  
--logging-volume cluster_6_log_vol`

**See also**

- ◆ [logging-volume add-mirror on page 292](#)
- ◆ [logging-volume create on page 293](#)

## ls

Displays information about the current object or context.

**Contexts** All contexts.

**Syntax**

```
ls options context
[-l|--long]
[-a|--attributes]
[-A|--no-attributes]
[-t|--attribute] selector
[-p|--paged]
[-m|--commands]
[-f|--full]
[-C|--no-contexts]
[-c|--context] context,context...
```

**Arguments** **Required arguments**

None.

**Optional arguments**

**[-l|--long]** - Display more detailed information.

**[-a|--attributes]** - Includes the attributes of the target context(s).

**[-A|--no-attributes]** - Excludes attributes.

**[-t|--attribute-selector] attribute-selector** - Display the contents of the specified attribute(s).

**[-p|--paged]** - Page output if it is longer than the window height.

**[-m|--commands] depth** - Includes commands in the listing.

**[-f|--full]** - Do not summarize long attribute listings.

**[-C|--no-contexts]** - Excludes contexts from the listing.

**[-c|--context] context,context...** - The context(s) whose contents are to be listed.

**Description** If the **[-c|--context]** argument is omitted, displays the contents of the current context.

The contents of a context include: its child contexts, if any; its attributes, if any; and the available commands, if any.

The context name can be any valid glob pattern.

**Notes** The VPLEX CLI includes **ll**, a pre-defined alias of 'ls -a'.

**Example** Display a device's attributes:

```
Vplexcli:/> ls -C /clusters/cluster-8/devices/device_CLAR0014_LUN04_1

/clusters/cluster-8/devices/device_CLAR0014_LUN04_1:
Name                               Value
-----
application-consistent            false
block-count                        2621440
block-size                          4K
capacity                           10G
geometry                           raid-0
health-indications                  []
health-state                        ok
locality                            local
operational-status                  ok
```

```

rebuild-allowed      -
rebuild-eta          -
rebuild-progress     -
.
.
.

```

Use the **--attribute** argument to display the operational status of all directors:

```
VPlexcli:/> ls --attribute /engines/*/directors/*::operational-status
```

```

/engines/engine-2-1/directors/dirB:
Name          Value
-----
operational-status  ok

```

```

/engines/engine-2-1/directors/dirA:
Name          Value
-----
operational-status  ok

```

```

.
.
.

```

Display a cluster's attributes and the contexts below the cluster context:

```
VPlexcli:/> ls /clusters/cluster-1
```

```
/clusters/cluster-1:
```

```

Attributes:
Name          Value
-----
allow-auto-join      true
auto-expel-count    0
auto-expel-period   0
auto-join-delay     0
cluster-id         1
connected           true
default-cache-mode  synchronous
default-caw-template true
director-names     [DirA, DirB]
island-id          1
operational-status  ok
transition-indications []
transition-progress []
health-state       ok
health-indications []

```

```

Contexts:
cluster-connectivity  devices          exports
  storage-elements
system-volumes        uninterruptible-power-supplies  virtual-volumes

```

Use a glob pattern to display all the fans:

```
VPlexcli:/> ls --long /**/fans
```

Use a glob pattern to display:

- ◆ All fans
- ◆ All the uninterruptable power supply settings:

```
VPlexcli:/> ls --long /**/fans, /**/uninterruptible-power-supplies/*
```

Use the **--attribute-selector** argument to display the contents of the 'virtual-volumes' attribute on all views:

```
Vplexcli:/> ls --attribute /clusters/*/exports/storage-views /*::virtual-volumes
```

**See also** ♦ [“alias” on page 27](#)

## management-server set-ip

Assigns IP address, net-mask, and gateway IP address to the specified management port.

**Contexts** All contexts.

**Syntax**

```
management-server set-ip
  [-i|--ip-netmask] destination IP address:subnet mask
  or
  [-i|--ip-netmask] destination IP address/CIDR mask
  [-g|--gateway] IP address
  [-p|--management-port] context path
```

**Arguments** **Required arguments**

**[-i|--ip-netmask] destination IP address:subnet mask** - The address and subnet mask of the Ethernet port. The format of the address/subnet mask depends on the version of IP.

- ◆ To specify an **IPv4** address - The format is: *destination IP address:subnet mask*  
For example: 172.16.2.0:255.255.255.0
- ◆ To specify an **IPv6** address - The format is: *destination IP address/CIDR netmask*  
For example: 3ffe:80c0:22c:803a:250:56ff:feb5:c1/64.

**[-g|--gateway] IP address** - The IP address of the gateway for this management server.

**[-p|--management-port] context path** - Ethernet port for which the parameter(s) are assigned/changed.

**Description** The VPLEX management server includes 4 Ethernet ports:

- ◆ eth0 - Service port.
- ◆ eth1 and eth2 - Internal management ports.
- ◆ eth3 - Public management port. The only Ethernet port in the VPLEX server that may be connected to an external management LAN.

The IP addresses for eth0, eth1, and eth2 cannot be changed.

Use the **management-server set-ip** command to set the IP address and netmask for port eth3.

Ports eth0, eth1 and eth2 do not have IPv6 addresses. Example output of the **ll** command for eth0, eth1 and eth2:

```
/management-server/ports/eth0:
Name          Value
-----
address       128.221.253.33
gateway       10.31.52.1
inet6-address []
inet6-gateway []
net-mask      255.255.255.224
```



### WARNING

*Changing the IP address for port eth3 can disrupt your inter-cluster link, and if VPLEX Witness is deployed, disrupt the VPN between the VPLEX clusters and the Witness Server.*

*Additional failures (for example a remote VPLEX cluster failure) while VPN between*

*the VPLEX clusters and the Witness Server is disrupted could lead to DU for all distributed virtual volumes in synchronous consistency groups.  
For the procedure to safely change the management server IP address, refer to the Troubleshoot => Cluster Witness section of the VPLEX Procedure Generator.*

**Example** Modify an IPv4 address:

- ◆ The **ll** command displays the current setting for eth3
- ◆ The **management-server set-ip** command modifies the port's IPv4 settings:

```
VPllexcli:/> ll /management-server/ports/eth3
```

| Name          | Values                  |
|---------------|-------------------------|
| address       | 10.31.52.207            |
| gateway       | 10.31.52.1              |
| inet6-address | 2001:f0d0:1002:11::2/64 |
| inet6-gateway | 2001:f0d0:1002:11::1    |
| net-mask      | 255.255.252.0           |

```
VPllexcli:/> management-server set-ip --ip-netmask  
10.31.52.197:255.255.252.0 --gateway 10.31.52.1 -p eth3
```

**Example** Modify an IPv6 address:

- ◆ The **ll** command displays the current setting for eth3
- ◆ The **management-server set-ip** command modifies the port's IPv6 settings:

```
VPllexcli:/> ll /management-server/ports/eth3
```

| Name          | Value                                                                                   |
|---------------|-----------------------------------------------------------------------------------------|
| address       | 10.31.52.197                                                                            |
| gateway       | 10.31.52.1                                                                              |
| inet6-address | [3ffe:80c0:22c:803c:215:17ff:fecc:4408/64,<br>3ffe:80c0:22c:803c:415:17ff:fecc:4408/64] |
| inet6-gateway | 3ffe:80c0:22c:803c::1                                                                   |
| net-mask      | 255.255.252.0                                                                           |

```
VPllexcli:/> management-server set-ip --ip-netmask  
3ffe:80c0:22c:803c:215:17ff:fed2:fe88/64 --gateway  
3ffe:80c0:22c:803c:1 -p eth3
```

**See also** ◆

---

## manifest upgrade

Loads a new manifest file, replacing the old one, if it exists.

**Contexts** All contexts.

**Syntax** `manifest upgrade  
[-m|--manifest] pathname`

**Arguments** **Required arguments**  
`[-m|--manifest] pathname` - Path to manifest file. Relative paths can be used.

**Description** The new manifest file will be validated before it replaces the old one.  
If there is no current valid manifest file (corrupted or missing), the specified manifest file is installed without confirmation.  
If a valid manifest file exists, confirmation is required if the specified manifest file does not have a newer version than the existing one.

**Example**

**See also** ♦ [“manifest version” on page 302](#)

---

## manifest version

Displays the version of the currently loaded manifest file.

**Contexts** All contexts.

**Syntax** `manifest version`

**Description** A Jython command used by scripts during upgrades.

**Example**  
VPLexcli:/> **manifest version**  
1.106.0

**See also** ♦ [“manifest upgrade” on page 301](#)

---

## meta-volume attach-mirror

Attaches a storage-volume as a mirror to a meta-volume.

**Contexts** All contexts.

**Syntax** meta-volume attach-mirror  
[-d|--storage-volume] *context path*  
[-v|--meta-volume] *context path*

**Arguments** **Required arguments**

**[-d|--storage-volume] *context path*** - Storage-volume to attach as a mirror to the meta-volume.

**[-v|--meta-volume] *context path*** - Meta-volume to which the storage-volume should be attached as a mirror.

**Description** Creates a mirror and backup of the specified meta-volume. The specified storage-volume(s):

- ◆ Must not be empty,
- ◆ Must be at the implied or specified cluster,
- ◆ Unclaimed,
- ◆ 78 GB or larger.

EMC recommends a mirror and a backup of the meta-volume are created using at least two disks from two different arrays.

---

**Note:** A mirror can be attached when the meta-volume is first created by specifying two storage volumes.

---

**Example** Attach storage-volume “VPD83T3:6...ade11” as a mirror to the existing meta-volume “\_dmx”:

```
Vplexcli:/engines/engine-1-1/directors> meta-volume attach-mirror --storage-volume  
VPD83T3:6006016023901d00484f496fa07ade11 --meta-volume _dmx
```

-volume \_dmx is created at /clusters/cluster-1/system-volumes.

**See also** ◆ [meta-volume detach-mirror on page 311](#)

## meta-volume backup

Creates a new meta-volume and writes the current in-memory system data to the new meta-volume without activating it.

**Contexts** All contexts.

**Syntax**

```
meta-volume backup
  [-d|--storage-volumes] context path,context path...
  [-c|--cluster] context path
  [-f|--force]
```

**Arguments** **Required arguments**

**[-d|--storage-volume] context path,context path...** \* List of two or more storage volumes to use in creating the backup meta-volume. The specified storage-volume(s):

- ◆ Must not be empty
- ◆ Must be at the implied or specified cluster
- ◆ Unclaimed
- ◆ 78 GB or larger

Type the system IDs for multiple (two or more) storage volumes, separated by commas.

**Optional arguments**

**[-c|--cluster] context path** - The cluster whose active meta-volume will be backed-up.

**[-f|--force]** - Forces the backup meta-volume to be activated without asking for confirmation.

\* - argument is positional.

**Description** Backup creates a point-in-time copy of the current in-memory metadata without activating it. The new meta-volume is named:

*current-metadata-namebackup\_YYYYMMMdd\_HHmms*

Metadata is read from the meta-volume only during the boot of each director.



### **IMPORTANT**

**At boot, VPLEX reads the first meta-volume it encounters. If an out-of-date meta-volume is visible to VPLEX, and is the first encountered during boot, that meta-volume is read.**

**All but the latest meta-volume should be moved to where they are not visible to VPLEX during boot.**

Create a backup meta-volume:

- ◆ As part of an overall system health check before a major migration or update
- ◆ If the VPLEX permanently loses access to both meta-volumes.



### **IMPORTANT**

**No modifications should be made to VPLEX during the backup procedure. Make sure that all other users are notified.**

---

Use the **II** command in the system-volumes context to verify that the meta-volume is **Active** and its **Ready** state is true.

**Example** Back up the metadata to a RAID 1 of two specified storage volumes:

```
Vplexcli:meta-volume backup -storage-volumes VPD83T3:60060480000190300487533030354636,  
VPD83T3:60060480000190300487533030343445
```

- See also**
- ◆ [meta-volume create on page 306](#)
  - ◆ [meta-volume destroy on page 310](#)

## meta-volume create

Creates a new meta-volume in a cluster when there is no existing active meta-volume.

**Contexts** All contexts.

**Syntax**

```
meta-volume create
  [-n|--name] name
  [-d|--storage-volumes] context path,context path...
  [-f|--force]
```

**Arguments** **Required arguments**

**[-n|--name] name** - \* Name of the new meta-volume.

**[-d|--storage-volume] context path,context path...** \* List of two or more storage volumes to use in creating the new meta-volume. The specified storage volumes must not be empty, and must be at the implied or specified cluster.

Type the system IDs for the storage volumes separated by commas.



### **IMPORTANT**

**Specify two or more storage volumes. Storage volumes should be on different arrays.**

### **Optional arguments**

**[f|--force]** - Forces the meta-volume to be created without asking for confirmation.

\* - argument is positional.

**Description**

VPLEX metadata includes virtual-to-physical mappings, data about devices, virtual volumes, and configuration settings.

Metadata is stored in cache and backed up on a specially designated external volume called the meta-volume.

The meta-volume is critical for system recovery. The best practice is to mirror the meta-volume across two or more back-end arrays to eliminate the possibility of data loss. Choose the arrays used to mirror the meta-volume such that they are not required to migrate at the same time.

Meta-volumes differ from standard storage volumes in that:

- ◆ A meta-volume is created without first being claimed,
- ◆ Meta-volumes are created directly on storage volumes, not extents.



### **CAUTION**

**If the meta-volume is configured on a CLARiiON array, it must not be placed on the vault drives of the CLARiiON.**

Performance is not critical for meta-volumes. The minimum performance allowed is 40 MB/sec and 100 4 K IOP/second.

The physical spindles for meta-volumes should be isolated from application workloads.

EMC recommends the following for meta-volumes:

- ◆ Read caching enabled.

- ◆ A “hot spare” meta-volume pre-configured in case of a catastrophic failure of the active meta-volume.
- ◆ Minimum of 78 GB.

If two or more storage-volumes are specified, they must be on two separate arrays if more than one array is present. This command creates a RAID-1 of all the storage-volumes.

**Example** In the following example:

- ◆ The **configuration show-meta-volume-candidates** command displays possible candidates:

**Note:** Example output is truncated. Vendor, IO Status, and Type fields are omitted.

- ◆ The **meta-volume create** command creates a new mirrored volume using the 2 specified storage volumes.
- ◆ The **ll** command displays the new meta-volume.

```

Vplexcli:/> configuration show-meta-volume-candidates
Name  Capacity...Array Name
-----
VPD83T3:60060480000190100547533030364539 187G .....EMC-SYMMETRIX-190100547
VPD83T3:60000970000192601707533031333132 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333133 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333134 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333135 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333136 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333137 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:60000970000192601707533031333138 98.5G.....EMC-SYMMETRIX-192601707
VPD83T3:6006016049e02100442c66c8890ee011 80G .....EMC-CLARiION-FNM00083800068
1
.
.
.
Vplexcli:meta-volume create --name ICO_META_1_1_Metadata -storage-volumes
VPD83T3:60000970000192601707533031333136, VPD83T3:60060480000190300487533030343445

Vplexcli:/>cd /clusters/cluster-1/system-volumes

Vplexcli:/clusters/cluster-1/system-volumes> ll ICO_META_1_1_Metadata

/clusters/cluster-1/system-volumes/ICO_META_1_1_Metadata:

Attributes:
Name  Value
-----
active  true
application-consistent  false
block-count   24511424
block-size  4K
capacity  79.5G
component-count   2
free-slots   31968
geometry  raid-1
health-indications                                       []
health-state   ok
locality   local
operational-status                                       ok
ready  true
rebuild-allowed  true
rebuild-eta   -
rebuild-progress   -
rebuild-status   done

```

```

rebuild-type          full
slots                 32000
stripe-depth          -
system-id             ICO_META_1_1_Metadata
transfer-size         2M
volume-type           meta-volume

```

Contexts:

| Name       | Description                                                               |
|------------|---------------------------------------------------------------------------|
| components | The list of components that support this device or system virtual volume. |

**Table 18 meta-volume display fields**

| Field                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| active                 | Indicates whether this is the currently-active metadata volume. The system has only one active metadata volume at a time.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| application-consistent | Whether or not this storage-volume is application-consistent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| block-count            | The number of blocks.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| capacity               | The size of the meta-volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| component-count        | The number of mirrors in this raid-1 meta-data volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| free-slots             | The number of free slots for storage-volume headers in this meta-volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| geometry               | Indicates the geometry or redundancy of this device. Will always be raid-1.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| health-indications     | If health-state is not "ok", additional information.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| health-state           | <p><b>ok</b> - The storage volume is functioning normally.</p> <p><b>degraded</b> - The storage volume may be out-of-date compared to its mirror. (This state applies only to a storage volume that is part of a RAID-1 Metadata Volume.)</p> <p><b>unknown</b> - VPLEX cannot determine the storage volume's Health state, or the state is invalid.</p> <p><b>non-recoverable error</b> - The storage volume may be out-of-date compared to its mirror (applies only to a storage volume that is part of a RAID-1 Metadata Volume), and/or VPLEX cannot determine the Health state.</p> <p><b>critical failure</b> - VPLEX has marked the storage volume as hardware-dead.</p> |
| locality               | <p>Locality of the supporting device.</p> <p><b>local</b> - The volume is local to the enclosing cluster.</p> <p><b>remote</b> - The volume is made available by a different cluster than the enclosing cluster, and is accessed remotely.</p> <p><b>distributed</b> - The virtual volume either has, or is capable of having, legs at more than one cluster. *}</p>                                                                                                                                                                                                                                                                                                            |
| operational status     | <p><b>ok</b> - The storage volume is functioning normally.</p> <p><b>degraded</b> - The storage volume may be out-of-date compared to its mirror. (This state applies only to a storage volume that is part of a RAID-1 Metadata Volume.)</p> <p><b>unknown</b> - VPLEX cannot determine the storage volume's Health state, or the state is invalid.</p> <p><b>error</b> - VPLEX has marked the storage volume as hardware-dead.</p> <p><b>starting</b> - The storage volume is not yet ready.</p> <p><b>lost-communication</b> - The storage volume is unreachable.</p>                                                                                                        |
| ready                  | Indicates whether this metadata volume is ready or not.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| rebuild-allowed        | Whether or not this device is allowed to rebuild.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

Table 18 meta-volume display fields

| Field            | Description                                                                                                                                                                                                                                                                                                                                                                              |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rebuild-eta      | The estimated time remaining for the current rebuild to complete.                                                                                                                                                                                                                                                                                                                        |
| rebuild-progress | The percentage of this device that has been rebuilt.                                                                                                                                                                                                                                                                                                                                     |
| rebuild-status   | The rebuild status of this device.                                                                                                                                                                                                                                                                                                                                                       |
| rebuild-type     | The rebuild type.<br><b>full</b> - A full copy of all the blocks.<br><b>incremental</b> - Uses a checksum differencing algorithm to transfer only those (chunks of) blocks that are different.<br><b>comparison</b> - A comparison copy.<br><b>resync</b> - A resync rewrites blocks that may have been affected by a director failure, guaranteeing that the mirror legs are identical. |
| slots            | The total number of slots for storage-volume headers in this meta-volume.                                                                                                                                                                                                                                                                                                                |
| stripe-depth     | The depth of a stripe in bytes when 'geometry' is 'raid-0'.                                                                                                                                                                                                                                                                                                                              |
| system-id        | Name assigned to the meta-volume.                                                                                                                                                                                                                                                                                                                                                        |
| transfer-size    | The transfer size during rebuild in bytes. See <a href="#">About transfer-size on page 53</a> .                                                                                                                                                                                                                                                                                          |
| volume-type      | For meta-volumes, this is always 'meta-volume'.                                                                                                                                                                                                                                                                                                                                          |

**See also** ♦ [meta-volume destroy on page 310](#)

## meta-volume destroy

Destroys a meta-volume, and frees its storage volumes for other uses.

**Contexts** All contexts.

**Syntax** meta-volume destroy  
[-v|--meta-volume] *context path*  
[-f|--force]

**Arguments** **Required arguments**

**[-v|--meta-volume] context path**- \* Meta-volume to destroy.

**Optional arguments**

**[f|--force]** - Destroys the meta-volume without asking for confirmation (allows the command to be run from a non-interactive script). Allows the meta-volume to be destroyed, even if the meta-volume is in a failed state and unreachable by VPLEX.

\* - argument is positional.

**Description** The meta-volume cannot be destroyed if its **active** attribute is true.

**Example** In the following example:

- ◆ It displays that the target meta-volume has an active state of false.
- ◆ The **meta-volume destroy** command destroys the meta-volume:

```
Vplexcli:/clusters/cluster-1/system-volumes> ll meta1
```

```
/clusters/cluster-1/system-volumes/meta1:
```

```
Attributes:
```

| Name                   | Value    |
|------------------------|----------|
| -----                  | -----    |
| active                 | false    |
| application-consistent | false    |
| block-count            | 23592704 |
| .                      |          |
| .                      |          |
| .                      |          |

```
Vplexcli:/clusters/cluster-1/system-volumes> meta-volume destroy -v  
meta1
```

```
Meta-volume 'meta1' will be destroyed. Do you wish to continue?  
(Yes/No) y
```

**See also** ◆ [meta-volume create on page 306](#)

## meta-volume detach-mirror

Detaches a storage-volume/mirror from a meta-volume.

**Contexts** All contexts.

**Syntax** meta-volume detach-mirror  
[-d|--storage-volume *context path*  
[-v|--meta-volume] *context path*  
[-s|--slot] *slot number*  
[f|--force]  
--discard

**Arguments** **Required arguments**

**[-d|--storage-volume] *context path*** - Storage volume to detach as a mirror from the meta-volume.

**[-v|--meta-volume] *context path*** - \* The meta-volume from which the storage-volume/mirror should be detached.

**Optional arguments**

**[-f|--force]** - Force the mirror to be discarded. Required when the **--discard** argument is used.

**[-s|--slot]** - The slot number of the mirror to be discarded. Applicable only when the **--discard** argument is used.

**--discard** - Discards the data in the mirror to be detached.

\* - argument is positional.

**Description** Detaches the specified storage volume from a meta-volume.

Use the **ll** command in `/clusters/cluster/system-volumes/meta-volume/components` context to display the slot number when using the **discard** argument.

### Example

```
Vplexcli:/clusters/cluster-1/system-volumes/meta-vol-1/components> ll
Name                               Slot  Type           Operational  Health  Capacity
-----                               Number -----          Status      State   -----
VPD83T3:60000970000192601869533030373030  2     storage-volume ok           ok      128G
VPD83T3:60000970000194900497533030333338  1     storage-volume ok           ok      128G
```

```
Vplexcli:/clusters/cluster-1/system-volumes/meta-vol-1/components> meta-volume detach-mirror
--storage-volume VPD83T3:60000970000194900497533030333338 --meta-volume meta-vol-1
```

**See also** ♦ [meta-volume attach-mirror on page 303](#)

## meta-volume move

Writes the current in-memory system data to the specified target meta-volume, then activates it.

**Contexts** All contexts.

**Syntax** `meta-volume move [-t|--target-volume context path]`

**Arguments** **Required arguments**

`[-t|--target-volume] context path` - Storage volume to move metadata to. Target volume must be:

- ◆ Unclaimed
- ◆ 78 GB or larger

**Description** Writes the metadata to the specified meta-volume, and activates it. The specified meta-volume must already exist (it is not created automatically). See [meta-volume create on page 306](#).

**Example** To move a meta-volume from one storage volume to another:

1. Identify one or more storage volume(s) that are:
  - Unclaimed
  - 78 GB or larger
2. Use the **meta-volume create** command to create a new meta-volume. Specify the storage volume identified in [Step 1](#) as the target volume(s).

```
Vplexcli:/engines/engine-1-1/directors> meta-volume create --name meta_dmx --storage-volumes VPD83T3:6006016037202200966da1373865de11
```

3. Use the **meta-volume move** command to move the existing in-memory metadata to the new meta-volume.

```
Vplexcli:/engines/engine-1-1/directors> meta-volume move --target-volume meta_dmx
```

**See also**

- ◆ [meta-volume create on page 306](#)
- ◆ [meta-volume destroy on page 310](#)

## meta-volume verify-on-disk-consistency

Analyzes a meta-volume's committed (on-disk) header slots for consistency across all mirrors/components.

**Contexts** All contexts.

**Syntax**

```
meta-volume verify-on-disk-consistency
[-l|--log] log file
[-f|--first] offset
[-n|--number] number
[-c|--cluster] cluster
[-m|--meta-volume] meta-volume
--style {short|long|slow}
```

**Arguments** **Required arguments**

**[-c|--cluster] cluster** - The cluster at which to analyze the active meta-volume. This argument may be omitted if the **--meta-volume** argument is present.

**[-m|--meta-volume] meta-volume** - The meta-volume to analyze. This argument may be omitted if the **--cluster** argument is present.

**[-l|--log] log file** - Full path to the log file on the management server.

**[-f|--first] first** - Offset of first header to analyze.

**[-n|--number] number** - Number of headers to analyze.

**--style {short | long | slow}** - The style of analysis to do. Valid values:

**short** - Requires special firmware support available only in “[notifications call-home view-event-modifications](#)” Release 5.0 and later.

**long** - Requires special firmware support available only in Release 5.0 and later.

**slow** - Available for all Release versions. Downloads the meta-volume headers from the meta-volume legs one at a time and compares them.



### **CAUTION**

**The slow option may take hours to complete on a production meta-volume.**

**Description**

An active meta-volume with an inconsistent on-disk state can lead to a data unavailability (DU) during NDU.

Best practice is to NDU immediately after passing this meta-volume consistency check.



### **IMPORTANT**

**If any errors are reported, do not proceed with the NDU, and contact EMC Customer Support.**

The format and the length of time for the command to complete vary depending on the VPLEX Release:

- ◆ For Release 4.2, the format of the command is:

```
meta-volume verify-on-disk-consistency -style slow --meta-volume meta-volume-name>
```

The length of time for the command to complete for Release 4.2 depends on the size of the configuration. On a very large configuration, the command may take as long as 4 hours.

---

**Note:** Running this command is optional before upgrading from Release 4.2.

---

- ◆ For Release 5.0 and later, the format of the command is:

```
meta-volume verify-on-disk-consistency -style long --meta-volume meta-volume-name>
```

The command takes 10-20 minutes to complete.

Check the report in the log file saved at: `/tmp/logfilename`. The log file reports mismatches between meta-volume RAID-1 legs.

If mismatches are detected, run the command again using the format:

```
meta-volume verify-on-disk-consistency -style slow --meta-volume meta-volume-name>
```

This version of the command takes an hour to complete.

---

**Note:** Running this command is recommended before upgrading from Release 5.0 or later.

---

**Example** Verify the specified meta-volume is consistent using the slow style:

```
Vplexcli:/> meta-volume verify-on-disk-consistency --style slow --meta-volume meta_cluster1
Doing a slow consistency check on meta-volume
'/clusters/cluster-1/system-volumes/meta_cluster1' for slots [0,32000).
Scanning offsets [0,32000)
.....
.....
.....
.....
.....
0 mismatches detected
```

Discover/display inconsistencies on a meta-volume using the long style:

```
Vplexcli:/clusters/cluster-2/system-volumes> meta-volume
verify-on-disk-consistency -c cluster-2 --style long
Doing a long consistency check on meta-volume
'/clusters/cluster-2/system-volumes/Cluster2_Meta_DGC_Vmax_mirror' for slots [0,32000).
Meta-volume is not consistent.
See /tmp/validatemeta.log for details of the inconsistencies.
```

**See also** ◆ [meta-volume create on page 306](#)

## monitor add-console-sink

Adds a console sink to the specified performance monitor.

**Contexts** All contexts.

In /monitoring context, command is **add-console-sink**.

**Syntax**

```
monitor add-console-sink
  [-o|--format] {csv|table}
  [-m|--monitor] monitor-name
  --force
```

**Arguments** **Required arguments**

**[-m|--monitor] *context path*** - \* Performance monitor to which to add a console sink.

**Optional arguments**

**[-f|--force]** - Forces the creation of the sink, even if existing monitors are delayed in their polling.

**[-o|--format] {csv|table}** - The output format. Can be csv (comma-separated values) or table.

Default: table.

\* -argument is positional.

**Description** Creates a console sink for the specified performance monitor. Console sinks send output to VPLEX Management Server Console.

Every monitor must have at least one sink, and may have multiple sinks. A monitor does not begin operation (polling and collecting performance data) until a sink is added to the monitor.

Use the **monitor add-console-sink** command to add a console sink to an existing monitor.



### **CAUTION**

**Console monitors display the specified statistics on the Unisphere for VPLEX, interrupting any other input/output to/from the console.**

**Example** Add a console sink with output formatted as table (the default output format for console sinks):

```
Vplexcli: /> monitor add-console-sink --monitor Director-2-1-B_TestMonitor
Navigate to the monitor context and use the ll console command to display the sink
settings:

Vplexcli: /cd /monitoring/directors/Director-2-1-B/monitors/Director-2-1-B_TestMonitor/sinks
Vplexcli: /monitoring/directors/Director-2-1-B/monitors/Director-2-1-B_TestMonitor/sinks> ll

Name      Enabled  Format  Sink-To
-----  -
console  true     table  console

Vplexcli: /monitoring/directors/Director-2-1-B/monitors/Director-2-1-B_TestMonitor/sinks> ll console

/monitoring/directors/Director-2-1-B/monitors/Director-2-1-B_TestMonitor/sinks/console:
Name      Value
-----  -
```

```
enabled true
format table
sink-to console
type console
```

- See also**
- ◆ [monitor add-file-sink on page 317](#)
  - ◆ [monitor remove-sink on page 324](#)
  - ◆ [monitor create on page 320](#)

## monitor add-file-sink

Adds a file sink to the specified performance monitor.

**Contexts** All contexts.

In /monitoring context, command is **add-file-sink**.

**Syntax**

```
monitor add-file-sink
  [-n|--name] name
  [-o|--format] {csv|table}
  [-m|--monitor] monitor-name
  [-f|--file] filename
  --force
```

**Arguments** **Required arguments**

**[-m|--monitor] context path** - \* Performance monitor to which to add a console sink.

**[-f|--file] filename** - \* File to which to send the sink's data.

**Optional arguments**

**[-f|--force]** - Forces the creation of the sink, even if existing monitors are delayed in their polling.

**[-n|--name] name** - Name for the new sink. If no name is provided, the default name "file" is applied.

**[-o|--format] {csv|table}** - The output format. Can be csv (comma-separated values) or table.

Default: csv.

\* -argument is positional.

**Description** Creates a file sink for the specified monitor. File sinks send output to the specified file.

The default location of the output file is /var/log/VPlex/cli.

The default name for the file sink context is 'file'.

Every monitor must have at least one sink, and may have multiple sinks. A monitor does not begin operation (polling and collecting performance data) until a sink is added to the monitor

Use the **monitor add-file-sink** command to add a file sink to an existing monitor.

**Example** To add a file sink to send output to the specified .csv file:

```
VPlexcli:/monitoring/directors/director-1-1-A/monitors> monitor add-file-sink --monitor
director-1-1-A_stats --file /var/log/VPlex/cli/director_1_1_A.csv
```

Navigate to the monitor sinks context and use the **ll sink-name** command to display the sink:

```
VPlexcli:/cd /monitoring/directors/director-1-1-A/monitors/director-1-1-A_stats/sinks
VPlexcli:/monitoring/directors/Director-1-1-A/monitors/director-1-1-A_stats/sinks> ll file

/monitoring/directors/Director-1-1-A/monitors/director-1-1-A_stats/sinks/file:
Name      Value
-----
enabled   true
format    csv
sink-to   /var/log/VPlex/cli/director_1_1_A.csv
```

type file

- See also**
- ◆ [monitor add-console-sink on page 315](#)
  - ◆ [monitor collect on page 319](#)
  - ◆ [monitor remove-sink on page 324](#)
  - ◆ [report create-monitors on page 392](#)

---

## monitor collect

Force an immediate poll and collection of performance data without waiting for the automatic poll interval.

**Contexts** All contexts.  
In /monitoring context, command is **collect**.

**Syntax** `monitor collect`  
`[-m|--monitors] context path,context path...`

**Arguments** **Required arguments**  
`[-m|--monitor] context path,context path` - One or more performance monitor(s) to update immediately.

**Description** Polls and collects performance data from user-defined monitors. Monitors must have at least one enabled sink.

### Example

```
Vplexcli:/> monitor collect  
/monitoring/directors/director-2-1-B/monitors/director-2-1-B_TestMonitor
```

```
Vplexcli:/>  
Source:                director-2-1-B_TestMonitor  
Time:                  2010-07-01 10:05:55  
director.be-ops (counts/s):  
.  
.  
.
```

- ◆ [monitor create on page 320](#)
- ◆ [report poll-monitors on page 395](#)

## monitor create

Creates a performance monitor.

**Contexts** All contexts.

In /monitoring context, command is **create**.

**Syntax**

```
monitor create
  [-p|--period] collection-period
  [-n|--name] monitor-name
  [-d|--director] context-path,context-path...
  [-s|--stats] stat,stat,stat...
  [-t|--targets] context-path,context-path...
  [-f|--force]
```

**Arguments** **Required arguments**

**[-n|--name] *monitor-name*** - \* Name of the monitor. The name is appended to the director on which the monitor is configured.

**[-s|--stats] *stat,stat,stat*** - \* One or more statistics to monitor, separated by commas.

Use the “[monitor stat-list](#)” command to display the available statistics.

**Optional arguments**

**[-p|--period] *collection-period***- Frequency at which this monitor collects statistics. Valid arguments are an integer followed by:

**ms** - milliseconds (period is truncated to the nearest second)

**s** - seconds (Default)

**min** - minutes

**h** - hours

**0** - Disables automatic polling.

The default period is 30 seconds.

**[-d|--director] *context path, context path...*** - \* List of one or more coma-separated director(s) for which to display statistics.

**[-t|--targets] *context path, context path...*** - List of one or more coma-separated targets for which to display statistics. Applicable only to statistics that require a target.

**[-f|--force]** - Forces the creation of the monitor, even if existing monitors are delayed in their polling.

\* - argument is positional.

**Description** Performance monitoring collects and displays statistics to determine how a port or volume is being used, how much I/O is being processed, CPU usage, and so on.

The VPLEX collects and displays performance statistics using two user-defined objects:

- ◆ **monitors** - Gather the specified statistics.
- ◆ **monitor sinks** - Direct the output to the desired destination. Monitor sinks include the console, a file, or a combination of the two.

The monitor defines the automatic polling period, the statistics to be collected, and the output of the format. The monitor sinks define the output destination.

Polling occurs when:

- ◆ The timer defined by the monitor's period attribute has expired.
- ◆ The monitor has at least one sink with the enabled attribute set to true.

Polling is suspended when:

- ◆ The monitor's period is set to 0, and/or
- ◆ All the monitor's sinks are either removed or their enabled attribute is set to false

Create short-term monitors to diagnose an immediate problem.

Create longer-term monitors for ongoing system management.

## About file rotation and timestamps

Starting in Release 5.1, the log files created by a monitor's file sink are automatically rotated when they reach a size of 10 MB. The 10MB file is saved as *filename.csv.n* where *n* is a number 1 - 10, and output is saved in a new file named *filename.csv.n+1*.

The .csv files are rotated up to 10 times.

In the following example, a monitor has exceeded 10MB of output. The initial 10MB are stored in *filename.csv.1*. Subsequent output is stored in *filename.csv*.

```
service@sms-cluster-1:/var/log/VPlex/cli> ll my-data.csv*
-rw-r--r-- 1 service users 2910722 2012-03-06 21:23 my-data.csv
-rw-r--r-- 1 service users 10566670 2012-03-06 21:10 my-data.csv.1
```

If the second file exceeds, 10B, it is saved as *filename.csv.2*, and subsequent output is saved in *filename.csv*. Up to 10 such rotations, and numbered .csv files are supported.

When the file sink is removed or the monitor is destroyed, output to the .csv file stops, and the current .csv file is timestamped. For example:

```
service@sms-cluster-1:/var/log/VPlex/cli> ll my-data.csv*
-rw-r--r-- 1 service users 10566670 2012-03-06 21:23 my-data.csv.1
-rw-r--r-- 1 service users 5637498 2012-03-06 21:26
my-data.csv_20120306092614973
```

## Procedure overview

To create and operate a monitor, use the following general steps:

1. Determine the type of statistic to collect from the target object.
 

Use the **monitor stat-list category** and/or the **monitor stat-list \*** command to display the statistics to include in the monitor.

Note whether the statistic you want to collect requires an argument (port number, volume ID).
2. Determine how often the monitor should collect statistics.
3. Use the **monitor create** command to create a monitor.
4. Use the **monitor add-sink** commands to add one or more sinks to the monitor.
 

Add a console sink to send performance data to the Unisphere for VPLEX.

Add a file sink to send performance data to a specified file.
5. The monitor begins operation (polling and collecting performance data) when the sink is added to the monitor.

To disable automatic polling without deleting the monitor or its sink(s), do one of the following:

- Use the **set** command to change the monitor's **period** attribute to 0.

- Use the **set** command to change the sink's **enabled** attribute to false.
6. Use the **monitor collect** command to update/collect statistics immediately without waiting for the monitor's automatic collection.
  7. Monitor output.

Console sinks display monitor output on the console.

For file sinks, navigate to `/var/log/VPlex/cli/` on the management server and use the **tail -f filename** to display the output,

or:

Send output to a csv file, open the file in Microsoft Excel and create a chart.

**Example** Create a simple monitor with the default period, and no targets:

```
VPlexcli:/monitoring> monitor create --name TestMonitor --director Director-2-1-B --stats
director.fe-read,director.fe-write
Successfully created 1 monitor(s) out of 1.
```

To create a monitor to collect statistics from the director category on `/engines/engine1/directors/Director-2-1-B` every 10 seconds:

```
VPlexcli:/monitoring> monitor create --name DirStats --period 10s --director
/engines/engine1/directors/Director-2-1-B --stats director.*
```

Create a monitor to collect statistics on all storage volumes at cluster-1:

```
VPlexcli:/monitoring> monitor create --name SVStats-Cluster1 --director
/engines/engine1/directors/Director-2-1-B --stats storage-volume.* --targets
/clusters/cluster-1/storage-elements/storage-volumes/*
```

Create a performance monitor to collect statistics on front-end port A0-FC01:

```
VPlexcli:/monitoring> monitor create --name FE-A0-FC01-stats --director
/engines/engine1/directors/Director-2-1-B --stats fe-prt.* --targets
/engines/engine1/directors/Director-2-1-B/hardware/ports/A0-FC01
```

- See also**
- ◆ [monitor add-console-sink on page 315](#)
  - ◆ [monitor add-file-sink on page 317](#)
  - ◆ [monitor destroy on page 323](#)
  - ◆ [monitor stat-list on page 325](#)
  - ◆ [report create-monitors on page 392](#)

## monitor destroy

Destroys a performance monitor.

**Contexts** All contexts.

In /monitoring context, command is **destroy**.

**Syntax**

```
monitor destroy
  [-m|--monitor] monitor-name,monitor-name...
  [-c|--context-only]
  [-f|--force]
```

**Arguments** **Required arguments**

**[-m|--monitor]** *monitor-name* - \* List of one or more names of the monitor(s) to destroy.

**Optional arguments**

**[-f|--force]** - Destroy monitors with enabled sinks and bypass confirmation.

**[-c|--context-only]** - Removes monitor contexts from the Unisphere for VPLEX and the CLI, but does not delete monitors from the firmware. Use this argument to remove contexts that were created on directors to which the element manager is no longer connected.

**Description** Deletes the specified performance monitor.

### Example

```
Vplexcli:/> monitor destroy Cluster_2_Dir_2B_diskReportMonitor,
Cluster_2_Dir_2B_portReportMonitor,Cluster_2_Dir_2B_volumeReportMonitor
WARNING: The following items will be destroyed:
```

Context

```
-----
/monitoring/directors/Cluster_2_Dir_2B/monitors/Cluster_2_Dir_2B_diskReportMonitor
/monitoring/directors/Cluster_2_Dir_2B/monitors/Cluster_2_Dir_2B_portReportMonitor
/monitoring/directors/Cluster_2_Dir_2B/monitors/Cluster_2_Dir_2B_volumeReportMonitor
```

Do you wish to proceed? (Yes/No) **y**

Monitor 'Cluster\_2\_Dir\_2B\_volumeReportMonitor' is owned by another management console and/or has enabled sinks. Do you wish to proceed ? (Yes/No) **y**

Monitor 'Cluster\_2\_Dir\_2B\_portReportMonitor' is owned by another management console and/or has enabled sinks. Do you wish to proceed ? (Yes/No) **y**

Monitor 'Cluster\_2\_Dir\_2B\_diskReportMonitor' is owned by another management console and/or has enabled sinks. Do you wish to proceed ? (Yes/No) **y**

**See also**

- ◆ [monitor create on page 320](#)
- ◆ [report create-monitors on page 392](#)

## monitor remove-sink

Removes a sink from a performance monitor.

**Contexts** All contexts.

In /monitoring context, command is **remove-sink**.

**Syntax** `monitor remove-sink  
[-s|--sinks] context path,context path...`

**Arguments** **Required arguments**

`[-s|--sinks] context path,context path...` - \* List of one or more sink(s) to remove.  
Entries must be separated by commas.

\* - argument is positional.

**Description** Removes one or more performance monitor sinks.

**Example** Remove a console sink:

```
Vplexcli:/monitoring/directors/director-2-1-B/monitors/director-2-1-B _TestMonitor> monitor  
remove-sink console
```

- ◆ [monitor add-console-sink on page 315](#)
- ◆ [monitor add-file-sink on page 317](#)

## monitor stat-list

Displays statistics available for performance monitoring.

**Contexts** All contexts.

In /monitoring context, command is **stat-list**.

**Syntax** monitor stat-list  
[-c|--categories] *category,category...*

**Arguments** **Required arguments**  
None.

### Optional arguments

**[-c|--categories] *category,category...*** - List of one or more statistics categories to display.

**Description** Performance statistics are grouped into categories Use the **monitor stat-list** command followed by the <Tab> key to display the statistics categories. For example:

```
Vplexcli:/monitoring> monitor stat-list
```

```
be-prt, cache, director, directory, fc-com-port, fe-director, fe-lu,  
fe-prt, ip-com-port, ramf, rdma, storage-volume, virtual-volume,  
wrt-pacing
```

Use the **--categories *categories*** argument to display the statistics available in the specified category. For example:

```
Vplexcli:/monitoring> monitor stat-list --categories director
Name                               Target  Type      Units
-----
director.be-aborts                 n/a    counter  counts/s
director.be-ops                    n/a    counter  counts/s
director.be-ops-read               n/a    counter  counts/s
director.be-ops-write              n/a    counter  counts/s
director.be-read                   n/a    counter  KB/s
director.be-write                  n/a    counter  KB/s
director.busy                      n/a    reading  %
director.fe-ops                   n/a    counter  counts/s
director.fe-ops-act                n/a    reading  counts
director.fe-ops-q                  n/a    reading  counts
director.fe-ops-read               n/a    counter  counts/s
director.fe-ops-write              n/a    counter  counts/s
director.fe-read                   n/a    counter  KB/s
director.fe-write                  n/a    counter  KB/s
director.heap-used                 n/a    reading  %
director.tcp-recv                  n/a    counter  KB/s
director.tcp-send                  n/a    counter  KB/s
```

Use the \* wildcard to display all statistics for all categories.

For example:

```
Vplexcli:/> monitor stat-list *
Name                               Target  Type      Units
-----
be-prt.read                        backend-port  counter  KB/s
be-prt.write                       backend-port  counter  KB/s
cache.dirty                        n/a        reading  KB
cache.miss                         n/a        counter  counts/s
cache.rhit                         n/a        counter  counts/s
cache.subpg                        n/a        counter  counts/s
director.be-aborts                 n/a        counter  counts/s
```

```
director.be-ops          n/a          counter counts/s
director.be-ops-read     n/a          counter counts/s
director.be-ops-write    n/a          counter counts/s
director.be-read         n/a          counter KB/s
director.be-write        n/a          counter KB/s
director.busy            n/a          reading  %
.
.
.
```

**See also** ♦ [monitor create on page 320](#)

## ndu pre-config-upgrade

Disruptively upgrades a VPLEX Geo that has not been fully installed and configured.

**Contexts** All contexts.

**Syntax**  
`ndu pre-config-upgrade`  
`[-u|--firmware] firmware tar file`  
`[-i|--image] firmware image file`

**Arguments** `[-u|--firmware] firmware tar file` - Full path to director firmware package on the management server. For example:

```
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.00.06-  
director-firmware-package.tar
```

`[-i|--image] firmware image file` - Full path to director firmware image on the management server. For example:

```
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.00.06-  
director-field-disk-image.tar
```

**Description** Disruptively upgrades a VPLEX Geo when the VPLEX is not fully installed and configured.



### CAUTION

**This procedure requires the VPLEX be in a pre-config state. Specifically, do not use this procedure unless NO meta-volume is configured (or discoverable).**

The following is a general outline of the procedure required to use the **ndu pre-config upgrade** command:

1. At the management server prompt, change the current directory to `/tmp/VPlexInstallPackages`.
2. Copy the current build's management server package:  
`VPlex-version-number>-management-server-package.tar.`
3. Install the new management server package using the VPlex-MS-Installer script with the `-m` option.

```
service@ManagementServer:/tmp/VPlexInstallPackages>
```

```
VPlex-MS-installer -m
```

```
VPlex-5.0.1.00.00.06-management-server-package.tar
```

```
Performing pre-install checks ok  
Clean up old configurations ok  
Install management server software ok  
Install base image patch ok  
Stop processes for Manufacturing Install ok  
Post manufacturing install configuration ok
```

4. Use the `sudo /sbin/shutdown -r now` command to reboot the management server.
5. Wait for the management server to come back up.
6. Copy the file `VPlex-5.0.1.00.00.06-director-field-disk-image.tar` to `/tmp/VPlexInstallPackages/`
7. At the management server prompt, type the `vpflexcli` command to connect to the VPLEX CLI.

8. Log in with username service and password.
9. Connect to each director using the **connect** command. For example, in a medium configuration:

```
VPLexcli:/> connect -o 128.221.252.35 --secondary-host 128.221.253.35 --name director-1-1-A  
Connected to Plex firmware director-1-1-A.
```

```
VPLexcli:/> connect -o 128.221.252.36 --secondary-host 128.221.253.36 --name director-1-1-B  
Connected to Plex firmware director-1-1-B.
```

```
VPLexcli:/> connect -o 128.221.252.37 --secondary-host 128.221.253.37 --name director-1-2-A  
Connected to Plex firmware director-1-2-A.
```

```
VPLexcli:/> connect -o 128.221.252.38 --secondary-host 128.221.253.38 --name director-1-2-B  
Connected to Plex firmware director-1-2-B.
```

10. Use the **ndu-pre-config-upgrade** command to start the upgrade.

```
VPLexcli:/> ndu pre-config-upgrade --image  
/tmp/VPLexInstallPackages/VPLex-5.0.1.00.00.06-director-field-disk-image.tar
```

```
=====  
[Tue Jul 12 10:50:04 2011] Pre-Configuration Image Upgrade starting  
=====  
*** Cluster /clusters/cluster-2 not involved in NDU ***  
=====  
Performing NDU pre-checks  
=====  
Verify director communication status.. OK  
Verify management network redundancy.. OK  
Verify management network latency.. OK  
Verify time drift between directors and management server.. OK  
Verify firmware software version can be retrieved.. OK  
Verify sufficient disk space on the management server for ndu.. OK  
Verify if the director SSD is about to trip.. OK  
Verify bios quiet mode enabled on the directors.. OK  
Verify no metadata volume available before configuration.. OK  
  
=====  
No problems found during NDU pre-check  
=====  
Setting up PXE Server...  
PXE Server Setup Finished...  
  
Upgrading OS for director [director-1-1-A,director-1-1-B]  
Rebooting the directors: .DONE  
Waiting for netboot. This will take several minutes.  
The directors have netbooted and they are now being imaged: .....DONE  
Rebooting the directors: .DONE  
Waiting for management connectivity: .....DONE  
* verifying director filesystems: .DONE  
* starting director firmware: .DONE  
Upgrading the BIOS & POST on [director-1-1-A,director-1-1-B].  
Rebooting directors [/engines/engine-1-1/directors/director-1-1-B,  
/engines/engine-1-1/directors/director-1-1-A]  
* rebooting: .DONE  
* waiting for management connectivity: .....DONE  
  
=====  
[Tue Jul 12 11:13:13 2011] Pre-Configuration Upgrade Finished Successfully  
=====  
The output for 'ndu pre-config-upgrade' has been captured in  
/var/log/VPLex/cli/capture/ndu-pre-config-upgrade-session.txt
```

- 
11. Use the **version -a** command to verify that the directors are running the correct version
  12. Use the **disconnect \*** command to disconnect from all the directors.
  13. Proceed to configure the system using EZ Setup.

- See also**
- ◆ [ndu start on page 348](#)
  - ◆ [ndu recover on page 335](#)
  - ◆ [ndu status on page 357](#)
  - ◆ VPLEX Procedure Generator

## ndu pre-check

Performs a pre-NDU validation and check.

**Contexts** All contexts.

**Syntax** `ndu pre-check`

**Description** Pre-NDU validation and checks include:

- ◆ Front-end high availability
- ◆ Back-end high availability
- ◆ Storage view configuration
- ◆ Unreachable storage volumes
- ◆ Devices undergoing rebuild
- ◆ Unhealthy virtual volumes
- ◆ Directors without front-end ports
- ◆ Director communication status
- ◆ Management network redundancy
- ◆ Management network latency
- ◆ Time drift between directors and management server
- ◆ Directors have sufficient free disk space for NDU
- ◆ Consistent O/S on all directors
- ◆ Directors are commissioned
- ◆ Firmware software version can be retrieved
- ◆ Validate system configuration
- ◆ Cluster status is ok
- ◆ Distributed device settings (Metro/Geo only)
- ◆ Inter-cluster communication connectivity (Metro/Geo only)
- ◆ Bios quiet mode is disabled
- ◆ Metadata on-disk-consistency
- ◆ Metadata backup configuration
- ◆ SSD is about to trip
- ◆ Stuck I/O
- ◆ Valid metadata volume
- ◆ Cluster witness state (if deployed)
- ◆ No running remote virtual volumes

### **Pre-checks for Geo systems**

- ◆ All asynchronous virtual-volumes are in a consistency group
- ◆ Asynchronous consistency groups use active-cluster-wins detach rule
- ◆ Consistency group detach rules configured
- ◆ No active I/O on losing (passive) cluster
- ◆ Winning cluster distributed asynchronous virtual volume legs are up-to-date

- ◆ Cache-mode consistency for distributed virtual volumes
- ◆ Consistency group configuration and current I/O status, action required



**CAUTION**

**NDU pre-checks must be run within 24 hours before starting the NDU process.**



**CAUTION**

**NDU is not supported in a VPLEX Geo configuration if remote volumes are exported. NDU does not proceed to the next release until the remote volumes are converted to local volumes or distributed asynchronous volumes.**

Disclaimers for multipathing in **ndu pre-check** give time for user to validate hosts.

If **ndu pre-check** detects that C4LX image upgrade is needed and the configuration is a VPLEX Metro, a prompt for the service account password on the remote management server is displayed.

Refer to the VPLEX Procedure Generator for more detailed information about NDU pre-checks.

**Example** In the following example, a VPLEX is ready for NDU:

```
Vplexcli:/> ndu pre-check
```

```
Warning:
```

```
During the NDU process, multiple directors will be offline for a
portion of the time. This is non-disruptive but is dependent on
a host-based multipathing solution being installed, configured,
and operating on all connected hosts.
```

```
Analyzing system configuration: .DONE
```

```
=====
```

```
Performing NDU pre-checks
```

```
=====
```

```
Verify director communication status..          OK
Verify management network redundancy..         OK
Verify management network latency..            OK
Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip..   OK
Verify bios quiet mode enabled on the directors.. OK
Verify all directors are running the same O/S.. OK
Verify directors have been commissioned..       OK
Verify no unreachable or dead storage-volumes.. OK
Verify no unhealthy virtual-volumes..          OK
Verify distributed device settings..            OK
Verify no unhealthy storage views..             OK
Verify storage view configuration..             OK
Verify valid system configuration..             OK
Verify valid metadata volume..                 OK
Verify cluster status..                         OK
Verify and prepare directors for ndu..          OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify no stuck I/O..                           OK
Verify meta-volume backup configuration..       OK
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version..  OK
Verify cache-mode consistency for distributed virtual-volumes.. OK
Verify cluster witness state..                 OK
```

```
=====
```

No problems found during NDU pre-check.

=====  
The output for 'ndu pre-check' has been captured in /var/log/VPlex/cli/capture/ndupre-check-session.txt  
VPlexcli: />

In the following example, a VPLEX is ready for NDU Geo:

VPlexcli: /> **ndu pre-check**

Warning:

During the NDU process, multiple directors will be offline for a portion of the time. This is non-disruptive but is dependent on a host-based multipathing solution being installed, configured, and operating on all connected hosts.  
Analyzing system configuration: .DONE

Warning: GeoPlex NDU is required!

The detected system configuration indicates that a GeoPlex NDU is required:

\* 680 virtual volumes in asynchronous cache mode

A GeoPlex NDU temporarily forces all asynchronous virtual-volumes into synchronous cache-mode. This requires turning off the inter-cluster link and will result in data unavailability for consistency groups and remote virtual-volumes at the losing cluster.

The --force-geo option is required to initiate NDU.

=====  
Performing NDU pre-checks

=====  
Verify director communication status.. OK  
Verify management network redundancy.. OK  
Verify management network latency.. OK  
Verify time drift between directors and management server.. OK  
Verify firmware software version can be retrieved.. OK  
Verify sufficient disk space on the management server for ndu.. OK  
Verify if the director SSD is about to trip.. OK  
Verify bios quiet mode enabled on the directors.. OK  
Verify all directors are running the same O/S.. OK  
Verify directors have been commissioned.. OK  
Verify no unreachable or dead storage-volumes.. OK  
Verify no unhealthy virtual-volumes.. OK  
Verify distributed device settings.. OK  
Verify no unhealthy storage views.. OK  
Verify storage view configuration.. OK  
Verify valid system configuration.. OK  
Verify valid metadata volume.. OK  
Verify cluster status.. OK  
Verify and prepare directors for ndu.. OK  
Verify the response time of front-end switches.. OK  
Verify the response time of back-end switches.. OK  
Verify no stuck I/O.. OK  
Verify meta-volume backup configuration.. OK  
Verify inter-cluster communications connectivity.. OK  
Verify the remote management server version.. OK  
Verify cache-mode consistency for distributed virtual-volumes.. OK  
Verify actions required for NDU Geo.. WARNING  
Verify asynchronous consistency groups use active-cluster-wins detach rule... OK  
Verify there is no active I/O on losing cluster.. OK  
Verify presence of all asynchronous virtual-volumes in consistency group... OK  
Verify no running remote virtual volumes.. OK  
Verify winning cluster asynchronous virtual-volume legs up-to-date.. OK

=====  
Warnings (1 warnings found)

[WARNING]:

Based on the configuration and last-reported I/O status of used consistency groups, the following behavior may be observed during NDU Geo:

| Consistency Group | I/O status                               | Auto-resume | During NDU Geo                                         | After NDU Geo                                                                                                                                                              |
|-------------------|------------------------------------------|-------------|--------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Win2k3_vm         | cluster-1: passive<br>cluster-2: passive | true        | virtual-volumes will be unavailable at<br>all clusters | I/O will resume automatically at losing cluster(s).<br>This may result in a change of the view of data for hosts at the losing cluster(s) resulting in possible data-loss. |

Below is a summary of any actions that may be required before and after NDU Geo:

| Consistency Group | Actions Required Before NDU                                                                                                                   | Actions Required After NDU |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| cg1               | Initiate I/O at the cluster for which you wish to continue I/O with during the NDU, otherwise I/O will be suspended at all involved clusters. |                            |

Note: Issuing 'ndu pre-check' with the --verbose option may yield more details.

The output for 'ndu pre-check' has been captured in /var/log/VPlex/cli/capture/ndu-pre-check-session.txt

VPlexcli:/>

In the following example, a VPLEX is not ready for NDU:

```
VPlexcli:/> ndu pre-check
Warning:
During the NDU process, multiple directors will be offline for a
portion of the time. This is non-disruptive but is dependent on
a host-based multipathing solution being installed, configured,
and operating on all connected hosts.
Analyzing system configuration: .DONE
=====
Performing NDU pre-checks
=====
Verify director communication status.. OK
Verify management network redundancy.. OK
Verify management network latency.. OK
Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip.. OK
Verify bios quiet mode enabled on the directors.. OK
Verify all directors are running the same O/S.. OK
Verify directors have been commissioned.. OK
Verify no unreachable or dead storage-volumes.. OK
Verify no unhealthy virtual-volumes.. OK
Verify distributed device settings.. OK
Verify no unhealthy storage views.. OK
Verify storage view configuration.. OK
Verify valid system configuration.. OK
Verify valid metadata volume.. OK
Verify cluster status.. OK
Verify and prepare directors for ndu.. OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify no stuck I/O.. OK
Verify meta-volume backup configuration.. ERROR
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version.. OK
Verify cache-mode consistency for distributed virtual-volumes.. OK
Verify cluster witness state.. OK
=====
Errors (1 errors found)
=====
[ERROR]:
```

No meta-volume backup scheduled for cluster-1. Run the 'configuration metadata-backup' command on the cluster-1 management server.

```
=====  
Note: Issuing 'ndu pre-check' with the --verbose option may yield more details.  
=====  
The output for 'ndu pre-check' has been captured in  
/var/log/VPlex/cli/capture/ndu-pre-check-session.txt
```

VPlexcli:/>

- See also**
- ◆ [ndu start on page 348](#)
  - ◆ [ndu recover on page 335](#)
  - ◆ [ndu status on page 357](#)
  - ◆ [VPLEX Procedure Generator](#)

## ndu recover

Perform NDU recovery after a failed NDU attempt.

**Contexts** All contexts.

**Syntax** `ndu recover`

**Description** If the NDU failed before I/O is transferred from the second upgraders (running old software) to the first upgraders (running new software) then the first upgraders are rolled back to the old software.

If the NDU failed after I/O transfer, the directors are rolled forward to the new software.

If no recovery is needed, a message is displayed.

It is safe to run the **ndu recover command** multiple times.

Recommendation: Run the first iteration of the **ndu recover** command on the same management server session used to do the initial NDU.

this command does the following:

1. Attempts to reboot any dead directors if the peer director in the same engine is still reachable.
2. Determines if an NDU recovery is required.
3. Determines the version to which to recover.
4. Determines what steps are required to recover the system to the desired version.
5. Executes the recovery plan.

### Example

```
Vplexcli:/> ndu recover
===== [Mon May 3
21:08:19 2010] Verifying management connectivity to directors.
===== Verified
management connectivity to directors ['director-1-1-A', 'director-1-1-B', 'director-1-2-A',
'director-1-2-B']
===== [Mon May 3
21:08:20 2010] Verifying director software version.
===== Not all
directors are running the same firmware version. {'v11.1.4-0':
[/engines/engine-1-1/directors/director-1-1-B, /engines/engine-1-2/directors/director-1-2-B],
None: [/engines/engine-1-1/directors/director-1-1-A,
/engines/engine-1-2/directors/director-1-2-A]} All directors are running from the currently
activated bank. Not all directors have the same version in the active bank.
===== [Mon May 3
21:08:22 2010] Proceeding with NDU Recovery
.
.
.
```

In the following example, VPLEX determines that an NDU recovery is not needed:

```
Vplexcli:/> ndu recover
=====
[Mon May 3 15:40:25 2010] Verifying management connectivity to directors.
=====
Verified management connectivity to directors ['director-1-1-A', 'director-1-1-B',
'director-2-1-A', 'director-2-1-B']
=====
```

[Mon May 3 15:40:27 2010] Verifying director software version.

```
=====
All directors are running the same firmware version.
All directors are running from the currently activated bank.
All directors have the same version in the active bank.
```

```
=====
[Mon May 3 15:40:29 2010] NDU recovery not needed.
=====
```

```
=====
[Mon May 3 15:40:29 2010] System state summary
=====
```

The plex is operational at director software version 0.0.0.0.20100428000000e with directors director-1-1-A, director-1-1-B, director-2-1-A, director-2-1-B.

```
=====
The output for 'ndu recover' has been captured in
/var/log/VPLEX/cli/capture/ndu-recover-session.txt
```

- See also**
- ◆ [ndu pre-check on page 330](#)
  - ◆ [ndu start on page 348](#)
  - ◆ [ndu status on page 357](#)
  - ◆ VPLEX Procedure Generator

## ndu rolling-upgrade c4lx-to-sles

Performs a rolling upgrade of each director's operating system from C4LX to SLES11.

**Contexts** All contexts.

**Syntax**

```
ndu rolling-upgrade c4lx-to-sles
[-i|--image] path to director SLES 11 image tar file
[--skip-be-switch-check]
[--skip-cluster-status-check]
[--skip-confirmations]
[--skip-distributed-device-settings-check]
[--skip-fe-switch-check]
[--skip-group-be-checks]
[--skip-group-fe-checks]
[--skip-group-health-checks]
[--skip-group-config-checks]
[--skip-meta-volume-backup-check]
[--skip-storage-volumes-check]
[--skip-sysconfig-check]
[--skip-view-config-check]
[--skip-view-health-check]
[--skip-virtual-volumes-check]
[--skip-wan-com-check]
```

**Arguments** **Required arguments**  
**[--image] path to director SLES 11 image tar file** - Full path to director image file on the management server. For example:

```
/tmp/VPlexInstallPackages/VPlex-4.2.0.00.00.11-director-
field-disk-image.tar
```

**Optional arguments**

For details about the optional `--skip` arguments, refer to [“ndu start” on page 348](#).

**Description** Runs a rolling upgrade from C4LX to SLES11.

Upgrades the OS for all the local directors one at a time.

If the directors are already on SLES11, this command does nothing.

NDU operates normally after the SLES11 upgrade.

**Note:** On VPLEX Metro configurations, rolling image upgrade automatically runs concurrently on the remote management server. The amount of time to complete the rolling upgrade varies depending on the number of engines at each cluster:

A VPLEX Metro with 1 engine per cluster takes about 40 minutes.

A VPLEX Metro with 2 engines per cluster takes approximately 80 minutes.

A VPLEX Metro with 4 engines per cluster, takes approximately 160 minutes.

**Example** Sample output is truncated.

```
VPlexcli: /> ndu rolling-upgrade c4lx-to-sles -i
/tmp/VPlexInstallPackages/VPlex-4.2.0.00.00.11-director-field-disk-image.tar
Directors running C4LX requiring rolling O/S upgrade to SLES ['director-1-1-B',
'director-1-1-A', 'director-2-1-B', 'director-2-1-A']
=====
Performing NDU pre-checks
=====
Verify director communication status.. OK
Verify management network redundancy.. OK
Verify management network latency.. OK
```

```

Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip.. OK
Verify bios quiet mode enabled on the directors.. OK
Verify directors have been commissioned.. OK
Verify no unreachable or dead storage-volumes.. OK
Verify no unhealthy virtual-volumes.. OK
Verify distributed device settings.. OK
Verify no unhealthy storage views.. OK
Verify storage view configuration.. OK
Verify valid system configuration.. OK
Verify valid metadata volume.. OK
Verify cluster status.. OK
Verify and prepare directors for ndu.. OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify no stuck I/O.. OK
Verify meta-volume backup configuration.. OK
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version.. OK
Verify cache-mode consistency for distributed virtual-volumes.. OK
Verify cluster witness state.. OK
Verify asynchronous consistency groups use active-cluster-wins detach rule... OK
Verify there is no active I/O on losing cluster.. OK

```

```

=====
No problems found during NDU pre-check
=====

```

```

Preparing management server at cluster-2 (this may take several minutes): .....DONE

```

```

=====
[Thu Jan 13 16:27:21 2011] O/S upgrade [/engines/engine-2-1/directors/director-2-1-A,
/engines/engine-1-1/directors/director-1-1-A]
=====

```

```

Waiting for system to be stable before reimaging [/engines/engine-2-1/directors/director-2-1-A,
/engines/engine-1-1/directors/director-1-1-A] : .DONE

```

```

backup director configs: [/engines/engine-2-1/directors/director-2-1-A,
/engines/engine-1-1/directors/director-1-1-A]
.
.
.

```

```

[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011]
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] fdownload: xfer_mode=7
min=1 max=65535 size=262144
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] OK .....
..... 19% 72.9M 0s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 50K .....
..... 39% 90.0M 0s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 100K .....
..... 58% 105M 0s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 150K .....
..... 78% 48.5M 0s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 200K .....
..... 97% 41.8M 0s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 250K ..... 100%
11444G=0.004s
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011]
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011] 2011-01-13 23:31:59
(64.9 MB/s) - `M8SB2-30UC-EMC 118032709.C04-6997' saved [262144/262144]
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:32:00 2011]
.
.
.

```

```

[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:34:55 2011] 86

```

```

[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:35:03 2011] 5
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:05 2011] 92
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:35:13 2011] 11
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:15 2011] 99
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:16 2011] 100
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:16 2011] Imaging completed.
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:16 2011] Verifying content on
/dev/sda against director.Wildcat.1.2.81.0.0.img
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:35:23 2011] 17
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:26 2011] 5
[PXE-Image director-1-1-A (128.221.252.35)] [Thu Jan 13 16:35:33 2011] 23
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:35:36 2011] 11
.
.
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] Verify passed. MD5
digest for director.Wildcat.1.2.81.0.0.img: bec2a0a2966552215ae0c1122b38a7d1
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] Removing config:
/tftpboot/pxelinux.cfg/01-00-60-16-10-40-52
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] removing sles payload
directory: /srv/www/htdocs/director
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] Rebooting the director
[128.221.252.94] (This will take a few minutes)
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] Tearing down pxe boot
configuration...
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] No configs to remove
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] No payloads to remove
[PXE-Image director-2-1-A (128.221.252.67)] [Thu Jan 13 16:38:06 2011] Shutting down DHCP
server ..done
* rebooting: .DONE
* waiting for management connectivity: .....DONE
restore director configs: [/engines/engine-1-1/directors/director-1-1-A,
/engines/engine-2-1/directors/director-2-1-A]
* starting director firmware: ..DONE
Waiting for system to be stable after reimaging [/engines/engine-2-1/directors/director-2-1-A,
/engines/engine-1-1/directors/director-1-1-A] : ...DONE

=====
[Thu Jan 13 16:46:00 2011] Successfully upgraded the O/S on
[/engines/engine-2-1/directors/director-2-1-A, /engines/engine-1-1/directors/director-1-1-A]
=====

=====
[Thu Jan 13 16:46:03 2011] O/S upgrade [/engines/engine-2-1/directors/director-2-1-B,
/engines/engine-1-1/directors/director-1-1-B]
=====
Waiting for system to be stable before reimaging [/engines/engine-2-1/directors/director-2-1-B,
/engines/engine-1-1/directors/director-1-1-B] : .DONE
backup director configs: [/engines/engine-2-1/directors/director-2-1-B,
/engines/engine-1-1/directors/director-1-1-B]
pxe director image: /diag/builds/cardiff/Vplex-4.2.0.00.00.11-director-field-disk-image.tar
[PXE-Image director-1-1-B (128.221.252.36)] [Thu Jan 13 16:46:07 2011] Configuring imager to
use payload in /diag/builds/cardiff/Vplex-4.2.0.00.00.11-director-field-disk-image.tar to
reimage director-1-1-b
[PXE-Image director-1-1-B (128.221.252.36)] [Thu Jan 13 16:46:07 2011] Preparing
.
.
* rebooting: .DONE
* waiting for management connectivity: .....DONE
restore director configs: [/engines/engine-1-1/directors/director-1-1-B,
/engines/engine-2-1/directors/director-2-1-B]
* starting director firmware: ..DONE
Waiting for system to be stable after reimaging [/engines/engine-2-1/directors/director-2-1-B,
/engines/engine-1-1/directors/director-1-1-B] :
..128.221.252.36/cpu0/log:5988:W/"006016103fe80224-2":460:<0>2011/01/14 00:02:18.35: ??? ?

```

.DONE

```
=====
[Thu Jan 13 17:03:20 2011] Successfully upgraded the O/S on
[/engines/engine-2-1/directors/director-2-1-B, /engines/engine-1-1/directors/director-1-1-B]
=====
Cleaning up management server at cluster-2:
.128.221.253.36/cpu0/log:5988:W/"006016103fe80224-2":460:<0>2011/01/14 00:02:18.35: ??? ?
DONE
=====
The output for 'ndu c4lx-to-sles' has been captured in
/var/log/VPlex/cli/capture/ndu-c4lx-to-sles-session.txt

VPlexcli:/>
```

- See also**
- ◆ [ndu start on page 348](#)
  - ◆ VPLEX Procedure Generator

## ndu rolling-upgrade ssd-fw

Starts a rolling upgrade of SSD firmware on the directors.

**Contexts** All contexts.

**Syntax**

```
ndu rolling-upgrade ssd-fw
[-i|--image] path to firmware image file
[-t|--targets] targets,targets,...
[--force]
[--check-only]
[--dry-run]
[--skip-be-switch-check]
[--skip-cluster-status-check]
[--skip-confirmations]
[--skip-distributed-device-settings-check]
[--skip-fe-switch-check]
[--skip-group-be-checks]
[--skip-group-config-checks]
[--skip-group-fe-checks]
[--skip-group-health-checks]
[--skip-meta-volume-backup-check]
[--skip-remote-mgmt-version-check]
[--skip-storage-volumes-check]
[--skip-sysconfig-check]
[--skip-view-config-check]
[--skip-view-health-check]
[--skip-virtual-volumes-check]
[--skip-wan-com-check]
```

### Arguments Required arguments

**[-i|--image] *firmware image file*** - \* Full path to director firmware image on the management server. For example:

```
/tmp/VPLexInstallPackages/VPLex-5.0.1.00.00.06-director-field-disk-image.tar
```

### Optional arguments

**[-t|--targets] *targets,targets,...*** - List of directors to upgrade.

**--force** - Must be specified to ignore SSD firmware version checking. (To upgrade to the same or older firmware)

**--check-only** - Check which directors will have their SSD firmware upgraded, not upgrade the firmware.

**--dry-run** - Do not perform the ssd firmware upgrade but run the same procedure as an actual install (including netbooting the directors).

**--skip-be-switch-check** - Skips the NDU pre-check for unhealthy back-end switches.

**--skip-cluster-status-check** - Skip the NDU pre-check for cluster problems (missing directors, suspended exports, inter-cluster link failure).

**--skip-confirmations** - Skip any user confirmations normally required before proceeding when there are NDU pre-check warnings.

**--skip-distributed-device-settings-check** - Skips the NDU pre-check for distributed device settings (auto-resume set to true).

**--skip-fe-switch-check** - Skips the NDU pre-check for unhealthy front-end switches.

**--skip-group-be-checks** - Skip all NDU pre-checks related to back-end validation. This includes the system configuration validation and unreachable storage volumes pre-checks.

**--skip-group-config-checks** - Skip all NDU pre-checks related to system configuration. This includes the system configuration validation and director commission pre-checks.

**--skip-group-fe-checks** - Skip all NDU pre-checks related to front-end validation. This includes the unhealthy storage views and storage view configuration pre-checks.

**--skip-group-health-checks** - Skip all NDU pre-checks related to system health validation. This includes the system configuration validation, unhealthy virtual volumes, cluster status, and the inter-cluster communications connectivity pre-checks.

**--skip-meta-volume-backup-check**- Skips the check to verify that backups for the meta-data volumes at all clusters have been configured.

**--skip-remote-mgmt-version-check** - Skip the remote management server version check.

**--skip-storage-volumes-check** - Skip the NDU pre-check for unreachable storage volumes.

**--skip-sysconfig-check** - Skip the system configuration validation NDU pre-check and proceeds with NDU even if there are errors with cache replication, logging volume setup, back-end connectivity, and metadata volume health.

**--skip-view-config-check** - Skip the NDU pre-check for storage view configuration (front-end high availability). This option is required to pass the NDU pre-checks when operating a minimum configuration. For minimum configurations, front-end high-availability pre-checks must be performed manually.

**--skip-view-health-check** - Skip the NDU pre-check for unhealthy storage views.

**--skip-virtual-volumes-check** - Skip the NDU pre-check for unhealthy virtual volumes.

**--skip-wan-com-check** - Skip the inter-cluster communications connectivity NDU pre-check and proceeds with NDU even if there are errors specifically related to inter-cluster communications connectivity.



### CAUTION

**Skipping the WAN communications pre-check may increase the risk for NDU failure should the inter-cluster communication connection fail.**

**Note:** Multiple skip options can be specified to skip multiple pre-checks. Enter skip options separated by a space.

### Description

Upgrades the directors one at a time. Assures that there are directors available to service I/O as some of the directors are being upgraded. The upgraded director rejoins the system before the next director is upgraded.

The director SSD firmware upgrade is performed by netbooting the director to ensure that the SSD is not in use while the firmware is being upgraded.

Non-disruptively upgrades the SSD firmware on the directors in a running VPLEX system.

Use this procedure for VPLEX 4.2 systems running SLES11 that did not have the SSD firmware upgrade incorporated in the **"[ndu rolling-upgrade c41x-to-sles](#)"** command.

## Procedure

1. Check to see which directors would be upgraded.

```
Vplexcli:/> ndu rolling-upgrade ssd-fw --check-only --image  
/tmp/VplexInstallPackages/Vplex-5.0.1.00.00.06-director-field-disk-image.tar
```

2. Perform the SSD firmware upgrade.

```
Vplexcli:/>ndu rolling-upgrade ssd-fw --image  
/tmp/VplexInstallPackages/Vplex-5.0.1.00.00.06-director-field-disk-image.tar
```

**Note:** This procedure takes approximately 10 minutes per director.

**Example:** Check whether the directors need the SSD firmware upgrade (--check-only option):

```
Vplexcli:/> ndu rolling-upgrade ssd-fw --image  
/tmp/VplexInstallPackages/Vplex-5.0.1.00.00.06-director-field-disk-image.tar --check-only
```

```
=====  
[Tue Feb 22 13:32:57 2011] Checking Director SSD Firmware  
=====  
[UPGRADE REQUIRED: director-1-1-B] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-1-1-A] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-2-1-A] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-2-1-B] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
=====  
The output for 'ndu ssd-fw' has been captured in  
/var/log/Vplex/cli/capture/ndu-ssd-fw-session.txt
```

Target specific directors for the SSD firmware upgrade using the --targets option:

```
Vplexcli:/> ndu rolling-upgrade ssd-fw --image /tmp/VplexInstallPackages/  
Vplex-5.0.1.00.00.06-director-field-disk-image.tar --targets director-*-2-A
```

```
=====  
[Mon Jan 24 15:45:26 2011] Checking Director SSD Firmware  
=====  
[SKIP: director-1-2-A] Current SSD firmware is already at [C08-6997]  
[SKIP: director-2-2-A] Current SSD firmware is already at [C08-6997]  
=====  
The output for 'ndu ssd-fw' has been captured in  
/tmp/derk/clidir/capture/ndu-ssd-fw-session.txt
```

Upgrade all the directors if the SSD firmware in the given image is newer:

```
Vplexcli:/> ndu rolling-upgrade ssd-fw --image  
/tmp/VplexInstallPackages/Vplex-5.0.1.00.00.06-director-field-disk-image.tar
```

```
=====  
[Tue Feb 22 13:32:57 2011] Checking Director SSD Firmware  
=====  
[UPGRADE REQUIRED: director-1-1-B] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-1-1-A] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-2-1-A] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
[UPGRADE REQUIRED: director-2-1-B] SSD Model:M8SB2-30UC-EMC 118032709, FwRev: C04-6693, To:  
C04-6997  
=====  
Performing NDU pre-checks  
=====  
Verify director communication status.. OK
```

```

Verify management network redundancy.. OK
Verify management network latency.. OK
Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip.. OK
Verify bios quiet mode enabled on the directors.. OK
Verify directors have been commissioned.. OK
Verify no unreachable or dead storage-volumes.. OK
Verify no unhealthy virtual-volumes.. OK
Verify distributed device settings.. OK
Verify no unhealthy storage views.. OK
Verify storage view configuration.. OK
Verify valid system configuration.. OK
Verify valid metadata volume.. OK
Verify cluster status.. OK
Verify and prepare directors for ndu.. OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify no stuck I/O.. OK
Verify meta-volume backup configuration.. OK
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version.. OK
Verify cache-mode consistency for distributed virtual-volumes.. OK
Verify cluster witness state.. OK
Verify asynchronous consistency groups use active-cluster-wins detach rule... OK
Verify there is no active I/O on losing cluster.. OK

```

```

=====
No problems found during NDU pre-check
=====

```

```

=====
[Tue Feb 22 13:48:08 2011] Setting up PXE Server on local management server
=====

```

```

=====
[Tue Feb 22 13:48:49 2011] SSD firmware upgrade /engines/engine-1-1/directors/director-1-1-B
=====

```

```

Waiting for system to be stable before upgrading the SSD firmware on
/engines/engine-1-1/directors/director-1-1-B : ...DONE
[PXE] Whitelisting director-1-1-b
Rebooting director-1-1-B: .DONE
Waiting for netboot. This will take several minutes.
upgrading the ssd firmware on /engines/engine-1-1/directors/director-1-1-B
SSD Desired FwRev: C04-6997
SSD Model: M8SB2-30UC-EMC 118032709, Current SSD FwRev: C04-6693
Upgrade SSD Firmware from C04-6693 to C04-6997
--2011-02-22 13:47:36-- http://128.221.253.33/director/M8SB2-30UC-EMC%20118032709.C04-6997
Connecting to 128.221.253.33:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 262144 (256K) [text/plain]
Saving to: `M8SB2-30UC-EMC 118032709.C04-6997'

```

```

    OK ..... 19% 11.7M 0s
   50K ..... 39% 11.1M 0s
  100K ..... 58% 10.9M 0s
  150K ..... 78% 11.1M 0s
  200K ..... 97% 11.2M 0s
 250K ..... 100% 11.7M=0.02s

```

```

2011-02-22 13:47:36 (11.2 MB/s) - `M8SB2-30UC-EMC 118032709.C04-6997' saved [262144/262144]

```

```

Downloaded M8SB2-30UC-EMC 118032709.C04-6997

```

```

/dev/sda:
Done.

fwdownload: xfer_mode=7 min=1 max=65535 size=262144

[PXE] Blacklisting director-1-1-b
Rebooting the netbooted director /engines/engine-1-1/directors/director-1-1-B: .DONE
Waiting for management connectivity: .....DONE
* verifying director filesystems: .DONE
* starting director firmware: .DONE

=====
[Tue Feb 22 14:00:02 2011] Successfully upgraded the SSD firmware on
/engines/engine-1-1/directors/director-1-1-B
=====

=====
[Tue Feb 22 14:00:02 2011] SSD firmware upgrade /engines/engine-1-1/directors/director-1-1-A
=====

...
...
...

=====
[Wed Feb 23 09:34:13 2011] Teardown PXE Server
=====
Waiting for system to be stable after upgrading the SSD firmware on directors: .....DONE
Cleaning up management server at cluster-2: .DONE

=====
[Wed Feb 23 09:36:22 2011] Director SSD Firmware upgrade summary
=====
director-1-1-B: SSD Model: M8SB2-30UC-EMC 118032709, Current FwRev: C04-6997, Image FwRev:
C04-6997
director-1-1-A: SSD Model: M8SB2-30UC-EMC 118032709, Current FwRev: C04-6997, Image FwRev:
C04-6997
director-2-1-B: SSD Model: M8SB2-30UC-EMC 118032709, Current FwRev: C04-6997, Image FwRev:
C04-6997
director-2-1-A: SSD Model: M8SB2-30UC-EMC 118032709, Current FwRev: C04-6997, Image FwRev:
C04-6997
=====
The output for 'ndu ssd-fw' has been captured in
/var/log/VPlex/cli/capture/ndu-ssd-fw-session.txt

```

Do a dry-run of the rolling upgrade of SSD firmware but skip the actual step of writing the new SSD firmware to the SSD device.

**Note:** Dry-run includes netbooting the directors.

```

VPlexcli:/> ndu rolling-upgrade ssd-fw --dry-run --image
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.06-director-field-disk-image.tar

```

Upgrade the SSD firmware on a subset of directors.

**--targets** selects only the A directors (one in each cluster in a VPlex Metro configuration)

**--force** proceeds even though the SSD firmware on the directors is already up-to-date

```

VPlexcli:/> ndu rolling-upgrade ssd-fw -i
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.06-director-field-disk-image.tar
--targets director-*-2-A --force

```

=====  
[Mon Jan 24 16:15:56 2011] Checking Director SSD Firmware  
=====

[UPGRADE REQUIRED: director-2-2-A] SSD Model:M8SB1-56UC-EMC 118032769, FwRev: C08-6997, To: C08-6997

[UPGRADE REQUIRED: director-1-2-A] SSD Model:M8SB1-56UC-EMC 118032769, FwRev: C08-6997, To: C08-6997  
=====

Performing NDU pre-checks  
=====

|                                                                 |    |
|-----------------------------------------------------------------|----|
| Verify director communication status..                          | OK |
| Verify management network redundancy..                          | OK |
| Verify management network latency..                             | OK |
| Verify time drift between directors and management server..     | OK |
| Verify directors have been commissioned..                       | OK |
| Verify firmware software version can be retrieved..             | OK |
| Verify no unreachable storage-volumes..                         | OK |
| Verify no unhealthy virtual-volumes..                           | OK |
| Verify distributed device settings..                            | OK |
| Verify no unhealthy storage views..                             | OK |
| Verify storage view configuration..                             | OK |
| Verify valid system configuration..                             | OK |
| Verify valid metadata volume..                                  | OK |
| Verify sufficient disk space on the management server for ndu.. | OK |
| Verify and prepare directors for ndu..                          | OK |

=====  
No problems found during NDU pre-check  
=====

Enter the 'service' userpassword for the cluster-2 management server:

=====  
[Mon Jan 24 16:17:18 2011] Upgrading SSD Firmware at cluster-1  
=====

=====  
[Mon Jan 24 16:17:18 2011] Setting up PXE Server on local management server  
=====

=====  
[Mon Jan 24 16:17:59 2011] SSD firmware upgrade /engines/engine-1-2/directors/director-1-2-A  
=====

Waiting for system to be stable before upgrading the SSD firmware on  
/engines/engine-1-2/directors/director-1-2-A : .DONE  
[PXE] Whitelisting director-1-2-a  
Rebooting director-1-2-A: .DONE  
Waiting for netboot. This will take several minutes.  
upgrading the ssd firmware on /engines/engine-1-2/directors/director-1-2-A  
SSD Desired FwRev: C08-6997  
SSD Model: M8SB1-56UC-EMC 118032769, Current SSD FwRev: C08-6997  
Skipping SSD Firmware Upgrade. [already at C08-6997]

.  
.  
.  
=====  
[Mon Jan 24 16:39:53 2011] Successfully upgraded the SSD firmware on  
/engines/engine-2-2/directors/director-2-2-A  
=====

=====  
[Mon Jan 24 16:39:53 2011] Teardown PXE Server  
=====

Waiting for system to be stable after upgrading the SSD firmware on directors: ...DONE  
Cleaning up management server at cluster-2: .DONE

```
=====
[Mon Jan 24 16:40:35 2011] Director SSD Firmware upgrade summary
=====
director-1-2-A: SSD Model: M8SB1-56UC-EMC 118032769, Current FwRev: C08-6997, Image FwRev:
C08-6997
director-2-2-A: SSD Model: M8SB1-56UC-EMC 118032769, Current FwRev: C08-6997, Image FwRev:
C08-6997
=====
The output for 'ndu ssd-fw' has been captured in
/tmp/derk/clidir/capture/ndu-ssd-fw-session.txt

VPlexcli:/>
```

- See also**
- ◆ [ndu start on page 348](#)
  - ◆ VPLEX Procedure Generator

## ndu start

Begins the non-disruptive upgrade (NDU) process of the director firmware.

**Contexts** All contexts.

**Syntax**

```
ndu start
[-u|--firmware] firmware tar file
[-p|--sleep-after-port-on] duration
[--cws-package] cws firmware tar file
[--force]
[--force-geo]
[--force-with-unreachable-cws]
[--skip-be-switch-check]
[--skip-cluster-status-check]
[--skip-confirmations]
[--skip-distributed-device-settings-check]
[--skip-fe-switch-check]
[--skip-group-be-checks]
[--skip-group-config-checks]
[--skip-group-fe-checks]
[--skip-group-health-checks]
[--skip-meta-volume-backup-check]
[--skip-storage-volumes-check]
[--skip-sysconfig-check]
[--skip-view-config-check]
[--skip-view-health-check]
[--skip-virtual-volumes-check]
[--skip-wan-com-check]
```

**Arguments** **Required arguments**

**[-u | --firmware] *firmware tar file*** - \* Full path to director firmware package on the management server. For example:

```
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.00.06-director-firmware-pack
age.tar
```

**Optional arguments**

**--cws-package *cws firmware tar file*** - Full path to Cluster Witness Server package on the management server. For example:

```
/tmp/VPlexInstallPackages/vplex-cws-5.0.0.00.12-upgrade.tar.gz
```

**Note:** Not required if upgrading to an official product release.

**[-p | -- sleep-after-port-on] *duration*** - Number of seconds VPLEX waits after FU front end ports are turned back on.

Default: 20 seconds.

**--force** - Ignore manifest checking of supported upgrades.

**--force-geo** - Force a VPLEX Geo NDU.



**CAUTION**

In order to start an NDU, the VPLEX must be fully installed and a meta-volume must be present. Do not run this command without first running the “**ndu pre-check**” command. For the director firmware upgrade to be non-disruptive to host I/O, the NDU pre-check must pass without any errors. The system must be healthy and in a highly available configuration.



### **CAUTION**

**--force-geo results in temporary data unavailability at the losing cluster for asynchronous consistency groups.**



### **CAUTION**

**Skipping the WAN communications pre-check may increase the risk for NDU failure should the inter-cluster communication connection fail.**

**--force-with-unreachable-cws** - Force the NDU to proceed if the Cluster Witness Server is unreachable and all clusters have administratively disabled the Cluster Witness component.

**--skip-be-switch-check** - Skips the NDU pre-check for unhealthy back-end switches.

**--skip-cluster-status-check** - Skip the NDU pre-check for cluster problems (missing directors, suspended exports, inter-cluster link failure, and so on).

**--skip-confirmations** - Skip any user confirmations normally required before proceeding when there are NDU pre-check warnings.

**--skip-distributed-device-settings-check** - Skips the NDU pre-check for distributed device settings (auto-resume set to true).

**--skip-fe-switch-check** - Skips the NDU pre-check for unhealthy front-end switches.

**--skip-group-be-checks** - Skip all NDU pre-checks related to back-end validation. This includes pre-checks for system configuration validation and unreachable storage volumes.

**--skip-group-config-checks** - Skip all NDU pre-checks related to system configuration. This includes the system configuration validation and director commission pre-checks.

**--skip-group-fe-checks** - Skip all NDU pre-checks related to front-end validation. This includes the unhealthy storage views and storage view configuration pre-checks.

**--skip-group-health-checks** - Skip all NDU pre-checks related to system health validation. This includes the system configuration validation, unhealthy virtual volumes, cluster status, and the inter-cluster communications connectivity pre-checks.

**--skip-meta-volume-backup-check** - Skips the check to verify that backups for the meta-data volumes at all clusters have been configured.

**--skip-storage-volumes-check** - Skip the NDU pre-check for unreachable storage volumes.

**--skip-sysconfig-check** - Skip the system configuration validation NDU pre-check and proceed with NDU even if there are errors with cache replication, logging volume setup, back-end connectivity, and metadata volume health.

**--skip-view-config-check** - Skip the NDU pre-check for storage view configuration (front-end high availability). This option is required to pass the NDU pre-checks when operating a minimum configuration. For minimum configurations, front-end high-availability pre-checks must be performed manually.

**--skip-view-health-check** - Skip the NDU pre-check for unhealthy storage views.

**--skip-virtual-volumes-check** - Skip the NDU pre-check for unhealthy virtual volumes.

**--skip-wan-com-check** - Skip the inter-cluster communications connectivity NDU pre-check and proceeds with NDU even if there are errors specifically related to inter-cluster communications connectivity.

## Description

This command:

- ◆ Prepares for the NDU:
  - Verifies the host environment, VPLEX environment, and for Metro and Geo configurations, verifies connectivity between clusters.
  - Downloads the new software package to all directors and installs the package into the inactive partition.
- ◆ Disables the Cluster Witness
- ◆ Starts the NDU:
  - Changes the state of the directors to upgrading state
  - Disables call-home
  - Designates A directors as the first upgraders (first set of directors to be upgraded).
  - Designates B directors as the second upgraders (second set of directors to be upgraded).
  - In VPLEX Metro and Geo configurations, first and second upgraders are split across clusters.
  - Shuts down the first upgraders, and waits for the second upgraders to stabilize.
  - All I/O is directed to the second upgraders. Applications running on the hosts do not experience disruption.
  - Reboots the first upgraders into the new software. First upgraders do not service host I/O until the NDU transfer occurs.
- ◆ Transfers host I/O from the second to the first upgraders:
  - Waits for the first upgraders to stabilize their view of the back-end devices, and become able to service I/O.
  - Enables the first upgrader front-end ports (without enabling I/O) so that host drivers can detect paths to the first upgraders.
  - Shuts down the second upgraders, and transfer I/O to the first upgraders.
  - All I/O is handled by the first upgraders running the new software. Applications running on the hosts do not experience disruption.
- ◆ Completes the NDU:
  - Reboots the second upgraders into the new software.
  - Restores all production settings.
  - Re-enables Cluster Witness.
  - Verifies system stability.
  - Verifies that the upgrade is successful.

### NDU with --force-geo option

When the **--force-geo** option is used, NDU performs the following steps:

- ◆ Pre-check validates the I/O patterns at active and passive clusters
- ◆ Disables Cluster Witness
- ◆ Disables the inter-cluster link
- ◆ Disables all wan-com ports (either FC or GigE)
- ◆ Waits for detach rules to engage
- ◆ Installs firmware packages on each director on the inactive partition
- ◆ Puts the system into write through mode

- ◆ Activates the new firmware on 1st upgraders
- ◆ Restores the inter-cluster link
  - 1st upgraders are now in write-back mode
  - Link is down for 10 minutes due to director reboot
- ◆ Performs I/O transfer from 2nd upgraders
- ◆ Activates the new firmware on 2nd upgraders
- ◆ Enables Cluster Witness
- ◆ Assesses whether log rebuilds are completed

## NDU pre-checks and skip options

Run the **“`ndu pre-check`”** command before running the **“`ndu start`”** command. The pre-checks executed by the **ndu pre-check** command verify that the upgrade from the current software to the new software is supported, the configuration supports NDU, and the system state is ready (clusters and volumes are healthy).

You must resolve all issues disclosed by the **ndu pre-check** command before running the **ndu start** command.

Skip options enable **ndu start** to skip one or more NDU pre-checks. Skip options should be used only after fully understanding the problem reported by the pre-check to minimize the risk of data unavailability.

**Note:** Skip options may be combined to skip more than one pre-check. Multiple skip options must be separated by a space.

On VPLEX Geo configurations, ndu pre-checks validate detach rules on consistency-groups and verify that I/O is not running on the passive-cluster.

### Before you begin

Before starting the NDU, perform the following tasks to verify that the VPLEX, host environment, and connectivity between clusters (for Metro and Geo configurations) are ready for upgrade:

- ◆ Run the **ndu pre-check** command.
- ◆ Verify that all host multipath applications are enabled and operating correctly. Applications should be set as follows.
  - HP-UX: PVLlinks set to failover. Consult EMC Host Connectivity Guide for HP-UX
  - VMWare: Set to Fixed Multipathing. Consult EMC Host Connectivity Guide for VMware ESX Server
  - IBM AIX: Native MPIO set to Round Robin. Consult EMC Host Connectivity Guide for IBM
  - AIX Linux: MPIO set to Round Robin Load Balancing. Consult EMC Host Connectivity Guide for Linux
  - All platforms: Powerpath set to Adaptive
- ◆ Use the **“`health-check`”** command to verify the VPLEX’s general health.
- ◆ For VPLEX Metro and Geo configurations, use the **“`vpn status`”** command to confirm that the VPN tunnel is established and that all local and remote directors are reachable.

Refer to the VPLEX Procedure Generator.

**Example** Start an NDU in a VPLEX Metro configuration:

```

VPlexcli:/> ndu start -u
/tmp/VPlexInstallPackages/VPLEX-5.0.1.00.00.06-directorfirmware-package.tar
=====
[Mon Apr 25 20:45:18 2011] Preparing for NDU
=====
Analyzing system configuration: .DONE
Warning:
During the NDU process, multiple directors will be offline for a
portion of the time. This is non-disruptive but is dependent on
a host-based multipathing solution being installed, configured,
and operating on all connected hosts.
Verifying director software package: .DONE
Current Director Software Version: 1.2.83.0.0
Target Director Software Version: 2.1.40.48.0
Cluster: /clusters/cluster-2
1st upgraders ['director-2-1-A']
2nd upgraders ['director-2-1-B']
Cluster: /clusters/cluster-1
1st upgraders ['director-1-1-A']
2nd upgraders ['director-1-1-B']
=====
Performing NDU pre-checks
=====
Verify director communication status.. OK
Verify management network redundancy.. OK
Verify management network latency.. OK
Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip.. OK
Verify bios quiet mode enabled on the directors.. OK
Verify all directors are running the same O/S.. OK
Verify directors have been commissioned.. OK
Verify no unreachable storage-volumes.. OK
Verify no unhealthy virtual-volumes.. OK
Verify distributed device settings.. OK
Verify no unhealthy storage views.. OK
Verify storage view configuration.. OK
Verify valid system configuration.. OK
Verify valid metadata volume.. OK
Verify cluster status.. OK
Verify and prepare directors for ndu.. OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version.. OK

=====
No problems found during NDU pre-check
=====
Shutting down Management Console web interface: .DONE
Shutting down background task scheduler: .DONE
Gathering device information: .DONE
Resetting director NDU state: .DONE
Transferring package files to directors: ....DONE
Installing firmware package on directors: .....DONE
=====
[Mon Apr 25 20:50:27 2011] Starting NDU
=====
Waiting for system to stabilize: .DONE
Waiting to shutdown 1st upgraders: .DONE
Disabling call-home: .DONE
Flush outstanding I/O on first upgraders: .DONE
Shutting down 1st upgraders: .DONE
Waiting for 2nd upgraders to stabilize: ..DONE
Rebooting 1st upgraders into target version (may take > 7 minutes):

```

```

* rebooting: .DONE
* waiting for management connectivity: .....DONE
* verifying director filesystems: ..DONE
* starting director firmware: ..DONE
* waiting for director firmware to recognize all devices: ....DONE
Waiting for ready to drain on 2nd upgraders: .DONE
=====
[Mon Apr 25 20:59:07 2011] Transferring I/O from 2nd to 1st upgraders
=====
Waiting for front-end on 1st upgraders to stabilize: .DONE
Enabling front-end on 1st upgraders: ..DONE
Flush outstanding I/O on second upgraders: .DONE
Shutting down 2nd upgraders and transferring I/O to 1st upgraders: .DONE
=====
[Mon Apr 25 20:59:40 2011] Finishing NDU
=====
Rebooting 2nd upgraders into target version (may take > 7 minutes):
* rebooting: ..DONE
* waiting for management connectivity: .....DONE
* verifying director filesystems: ..DONE
* starting director firmware: ..DONE
* waiting for director firmware to recognize all devices: ...DONE
* waiting for front-end on 2nd upgraders to stabilize: .DONE
Checking system stability after NDU finished: ..DONE
=====
[Mon Apr 25 21:07:45 2011] NDU Finished Successfully
=====
Restoring Management Console web interface: .DONE
=====
[Mon Apr 25 21:07:51 2011] System state summary
=====
The directors {director-1-1-A, director-1-1-B, director-2-1-A, director-2-1-B}
are operational at version 2.1.40.48.0.
Directors already running SLES ['director-1-1-A', 'director-1-1-B', 'director-2-1-
A', 'director-2-1-B']
No reachable directors requiring [C4LX to SLES] O/S upgrade
=====
The output for 'ndu start' has been captured in /var/log/Vplex/cli/capture/ndustart-
session.txt
Vplexcli:>

```

**Example** Start an NDU in a VPLEX Geo configuration:

```

Vplexcli:> ndu start --force-geo -u
/tmp/VplexInstallPackages/Vplex-5.0.1.00.00.06-director-firmware-package.tar

[13:29:39] =====
[13:29:39] [Thu Jul 21 13:29:39 2011] Preparing for NDU
[13:29:39] =====
Analyzing system configuration: .DONE

Warning: GeoPlex NDU is required!
The detected system configuration indicates that a GeoPlex NDU is required:

* 680 virtual volumes in asynchronous cache mode

A GeoPlex NDU temporarily forces all asynchronous virtual-volumes into
synchronous cache-mode. This requires turning off the inter-cluster link and
will result in data unavailability for consistency groups and remote
virtual-volumes at the losing cluster.

Warning:
During the NDU process, multiple directors will be offline for a
portion of the time. This is non-disruptive but is dependent on
a host-based multipathing solution being installed, configured,

```

```

and operating on all connected hosts.
Verifying director software package: ..DONE
Current Director Software Version: 2.1.40.48.0
Target Director Software Version: 2.1.40.51.0
Cluster: /clusters/cluster-1
  1st upgraders ['director-1-1-A']
  2nd upgraders ['director-1-1-B']
Cluster: /clusters/cluster-2
  1st upgraders ['director-2-1-A']
  2nd upgraders ['director-2-1-B']

```

```

=====
Performing NDU pre-checks
=====

```

```

Verify director communication status.. OK
Verify management network redundancy.. OK
Verify management network latency.. OK
Verify time drift between directors and management server.. OK
Verify firmware software version can be retrieved.. OK
Verify sufficient disk space on the management server for ndu.. OK
Verify if the director SSD is about to trip.. OK
Verify bios quiet mode enabled on the directors.. OK
Verify all directors are running the same O/S.. OK
Verify directors have been commissioned.. OK
Verify no unreachable or dead storage-volumes.. OK
Verify no unhealthy virtual-volumes.. OK
Verify distributed device settings.. OK
Verify no unhealthy storage views.. OK
Verify storage view configuration.. OK
Verify valid system configuration.. OK
Verify valid metadata volume.. OK
Verify cluster status.. OK
Verify and prepare directors for ndu.. OK
Verify the response time of front-end switches.. OK
Verify the response time of back-end switches.. OK
Verify no stuck I/O.. OK
Verify meta-volume backup configuration.. OK
Verify inter-cluster communications connectivity.. OK
Verify the remote management server version.. OK
Verify cache-mode consistency for distributed virtual-volumes.. OK
Verify actions required for NDU Geo.. WARNING
Verify asynchronous consistency groups use active-cluster-wins detach rule.. OK
Verify there is no active I/O on losing cluster.. OK
Verify presence of all asynchronous virtual-volumes in consistency group.. OK
Verify no running remote virtual volumes.. OK
Verify winning cluster asynchronous virtual-volume legs up-to-date.. OK

```

```

=====
Warnings (1 warnings found)
=====

```

```

[WARNING]:
Based on the configuration and last-reported I/O status of used consistency groups, the
following
behavior may be observed during NDU Geo:

```

| Consistency Group | I/O status                               | Auto-resume | During NDU Geo                                      | After NDU Geo                                                                                                                                                           |
|-------------------|------------------------------------------|-------------|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cg1               | cluster-1: passive<br>cluster-2: passive | true        | virtual-volumes will be unavailable at all clusters | I/O will resume automatically at losing cluster(s). This may result in a change of the view of data for hosts at the losing cluster(s) resulting in possible data-loss. |
| cg1               | cluster-1: passive<br>cluster-2: passive | true        | virtual-volumes will be unavailable at all clusters | I/O will resume automatically at losing cluster(s). This may result in a change of the view of data for hosts at the losing                                             |

|     |                                                   |      |                                                                                                      |                                                                                                                                                                                                                                             |
|-----|---------------------------------------------------|------|------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cg3 | cluster-1: active<br>writes<br>cluster-2: passive | true | virtual-volume I/O<br>will continue at<br>cluster-1 and suspend<br>at all other involved<br>clusters | cluster(s) resulting in<br>possible data-loss.<br>I/O will resume<br>automatically at losing<br>cluster(s). This may result<br>in a change of the view of<br>data for hosts at the losing<br>cluster(s) resulting in<br>possible data-loss. |
|-----|---------------------------------------------------|------|------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Below is a summary of any actions that may be required before and after NDU Geo:

| Consistency Group | Actions Required Before NDU                                                                                                                   | Actions Required After NDU |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| -----             | -----                                                                                                                                         | -----                      |
| cg1               | Initiate I/O at the cluster for which you wish to continue I/O with during the NDU, otherwise I/O will be suspended at all involved clusters. |                            |
| cg2               | Initiate I/O at the cluster for which you wish to continue I/O with during the NDU, otherwise I/O will be suspended at all involved clusters. |                            |

Do you wish to continue anyway (Yes/No) **yes**

```

=====
Note: Issuing 'ndu pre-check' with the --verbose option may yield more details.
Shutting down Management Console web interface: .DONE
DONE
Shutting down background task scheduler: Forcing I/O suspension according to loser rules: .DONE
Waiting until losing clusters become passive: .DONE
Disabling the inter-cluster com link: .DONE
Validating detach rule engagement: .WARNING: NDU will wait for at least 5 seconds for all detach
rules to engage.
DONE
Gathering device information: .DONE
Resetting director NDU state: .DONE
Transferring package files to directors: ...DONE
Installing firmware package on directors: .DONE
[13:45:43] =====
[13:45:43] [Thu Jul 21 13:45:43 2011] Pre-Ndu: Director's Initiator Login Info
[13:45:43] =====
[13:45:44] =====
[13:45:44] [Thu Jul 21 13:45:44 2011] Starting NDU
[13:45:44] =====
Waiting for system to stabilize: .DONE
Waiting to shutdown 1st upgraders: .DONE
Disabling call-home: .DONE
Flush outstanding I/O on first upgraders: .DONE
Shutting down 1st upgraders: .DONE
Waiting for 2nd upgraders to stabilize: ..DONE
Rebooting 1st upgraders into target version (may take > 7 minutes):
* rebooting: .DONE
* waiting for management connectivity: .....DONE
* verifying director filesystems: .DONE
* starting director firmware: .DONE
* restoring 1st upgrader inter-cluster com link: .DONE
* waiting for director firmware to recognize all devices: .....DONE
Waiting for ready to drain on 2nd upgraders: .DONE
[13:53:33] =====
[13:53:33] [Thu Jul 21 13:53:33 2011] Transferring I/O from 2nd to 1st upgraders
[13:53:33] =====
Waiting for front-end on 1st upgraders to stabilize: .DONE
Enabling front-end on 1st upgraders: ..DONE
Flush outstanding I/O on second upgraders: .DONE
Shutting down 2nd upgraders and transferring I/O to 1st upgraders: .
[13:54:01] =====

```

```

[13:54:01] [Thu Jul 21 13:54:01 2011] IO-Transfer: Director's Initiator Login Info
[13:54:01] =====
DONE
Resume suspended I/O according to loser rules: .DONE
[13:54:05] =====
[13:54:05] [Thu Jul 21 13:54:05 2011] Finishing NDU
[13:54:05] =====
Rebooting 2nd upgraders into target version (may take > 7 minutes):
* rebooting: ..DONE
* waiting for management connectivity: .....DONE
* verifying director filesystems: .DONE
* starting director firmware: ..DONE
* restoring 2nd upgrader inter-cluster com link: .DONE
* waiting for director firmware to recognize all devices: ....DONE
* waiting for front-end on 2nd upgraders to stabilize: .DONE
Checking system stability after NDU finished: .DONE
[14:00:47] =====
[14:00:47] [Thu Jul 21 14:00:47 2011] NDU Finished Successfully
[14:00:47] =====
[14:00:48] =====
[14:00:48] [Thu Jul 21 14:00:48 2011] Post-Ndu: Director's Initiator Login Info
[14:00:48] =====
Restoring Management Console web interface: .DONE
Restoring background task scheduler: .DONE
[14:00:49] =====
[14:00:49] [Thu Jul 21 14:00:49 2011] System state summary
[14:00:49] =====
The directors {director-1-1-B, director-1-1-A, director-2-1-B, director-2-1-A} are
operational at version 2.1.40.51.0.
Directors already running SLES ['director-1-1-B', 'director-1-1-A', 'director-2-1-B',
'director-2-1-A']
No reachable directors requiring [C4LX to SLES] O/S upgrade

```

**Example** Start an NDU with multiple skip options:

```

VPlexcli:/> ndu start --firmware
/tmp/VPlexInstallPackages/VPlex-5.0.1.00.00.06-director-firmware-package.tar
--skip-storage-volumes-check --skip-virtual-volumes-check

```

- See also**
- ◆ [ndu pre-check on page 330](#)
  - ◆ [ndu recover on page 335](#)
  - ◆ [ndu status on page 357](#)
  - ◆ VPLEX Procedure Generator

## ndu status

Displays the NDU status.

**Contexts** All contexts.

**Syntax** ndu status [--verbose]

**Description** If an NDU firmware or OS upgrade is running, this command displays the upgrade activity.

If neither NDU firmware or OS upgrade is running, this command displays information about the previous NDU firmware upgrade.

If the last operation was a rolling-upgrade, the OS upgrade information is displayed. The **ndu start** command clears this information.

If an NDU firmware or OS upgrade has failed, this command displays a message to use the **ndu recover** command.

if an NDU recovery is in progress, has succeeded or failed, this command displays a status message.

**Example** Display a successful NDU after completion:

```
Vplexcli: /> ndu status
Gathering NDU status...
```

```
No firmware or OS upgrade in progress.
Last Firmware Upgrade attempt on Fri, 17 Dec 2010 01:07:51
  From version 2.1.17.0.0 to version 2.1.19.0.0
  Was started on management server 10.6.209.61
  Result: succeeded
```

```
=====
[Fri Dec 17 18:05:21 2010] System state summary
=====
  The directors {director-1-1-B, director-1-1-A, director-1-2-B, director-1-2-A} are
operational at version 2.1.19.0.0.
=====
The output for 'ndu status' has been captured in
/var/log/Vplex/cli/capture/ndu-status-session.txt
```

Display NDU status after an NDU failed and **ndu recover** was run:

```
Vplexcli: /> ndu status
Gathering NDU status...
```

```
No firmware or OS upgrade in progress.
Last Firmware Upgrade attempt on Fri, 17 Dec 2010 00:39:29
  From version 2.1.19.0.0 to version None
  Was started on management server 10.6.209.61
  Result: failed
  Reason: Encountered a problem while preparing to start the NDU.
Unable to extract director package files, return code 2.
```

```
NDU recover succeeded on management server 127.0.0.1 on Fri, 17 Dec 2010 01:00:27.
```

```
=====
[Fri Dec 17 01:05:25 2010] System state summary
=====
  The directors {director-1-1-B, director-1-1-A, director-1-2-B, director-1-2-A} are
operational at version 2.1.19.0.0.
=====
```

The output for 'ndu status' has been captured in  
`/var/log/VPlex/cli/capture/ndu-status-session.txt`

- See also**
- ◆ [ndu pre-check on page 330](#)
  - ◆ [ndu start on page 348](#)
  - ◆ [ndu recover on page 335](#)
  - ◆ [VPLEX Procedure Generator](#)

## notifications call-home import-event-modifications

Imports and applies modified call-home events.

**Contexts** All contexts.

In /notifications context, command is **call-home import-event-modifications**.

In /notifications/call-home context, command is **import-event-modifications**.

**Syntax** notifications call-home import-event-modifications  
[-m|--modified-events-file] *filename*  
[-f|--force]

**Arguments** **Required arguments**

**[-m|--modified-events-file] *file*** - Path to the file containing the modified call-home events.

**Optional arguments**

**[-f|--force]** - Forces the import of the specified file without asking for confirmation. Allows this command to be run from non-interactive scripts.

**Description**

Imports and applies modifications to call-home events. This command imports the specified .xml file that contains modified call-home events. There can be two types of .xml event files:

- ◆ **EMC-generic** events are modifications recommended by EMC.

EMC provides an .xml file containing commonly requested modifications to the default call-home events.

- ◆ **Customer-specific** events are events modified to meet a specific customer requirement.

EMC provides a custom events file developed by EMC engineering and applied by EMC Technical Support.

Call-home behaviors changes immediately when the modified events file is applied.

If a customized events file is already applied, applying a new file overrides the existing file.

If the same event is modified in both the customer-specific and EMC-generic files, the modification specified in the customer-specific file is applied for that event, and the note "Not applied" appears in the command output.

If call-home is disabled when the custom events file is applied, the modified events are saved and applied when call-home is enabled.

Use the **set** command to enable/disable call-home notifications. Refer to [Enable/disable call-home notifications on page 445](#).

Use the **ls notifications/call-home** command to display whether call-home is enabled:

```
Vplexcli:/> ls /notifications/call-home

/notifications/call-home:

Attributes:
Name      Value
-----  -
enabled   true
```

**Example** In the following example, the custom events file 'custom.xml' is imported from a directory on the management server and applied when call-home is disabled:

```
Vplexcli:/notifications/call-home> import-event-modifications -m  
/home/service/demo/custom.xml
```

Importing the 'custom\_events.xml' file will override the existing call-home events, for all the event categories specified in the 'custom\_events.xml' file.

Do you want to proceed? (Yes/No) **yes**

The events provided in '/home/service/Test/custom\_events.xml' is saved. These events will be applied when call-home is enabled.

```
Vplexcli:/notifications/call-home>
```

**Example** In the following example, a custom events file is imported when call-home notifications is enabled:

```
Vplexcli:/notifications/call-home> import-event-modifications --file  
/home/service/Test/customCallHome.xml
```

Importing the 'custom\_events.xml' file will override the existing call-home events, for all the event categories specified in the 'custom\_events.xml' file.

Do you want to proceed? (Yes/No) **yes**

- See also**
- ◆ [notifications call-home remove-event-modifications on page 361](#)
  - ◆ [notifications call-home view-event-modifications on page 363](#)
  - ◆ [notifications call-home test on page 365](#)

## notifications call-home remove-event-modifications

Removes customized call-home events files, including customer-specific modifications and modifications recommended by EMC.

**Contexts** All contexts.

In /notifications context, command is **call-home remove-event-modifications**.

In /notifications/call-home context, command is **remove-modifications**.

**Syntax** `notifications call-home remove-event-modifications`  
`[-c|--customer-specific]`  
`[-e|--emc-generic]`  
`[-f|--force]`

**Arguments** **Optional arguments**

**[-c|--customer-specific]** - If a customer-specific call-home events file has been imported, removes the file.

**[-e|--emc-generic]** - If an EMC call-home events file has been imported, removes the file.

**[-f|--force]** - Removes the specified imported call-home events file without asking for confirmation. Allows this command to be executed from a non-interactive script.

**Description** This command removes the specified custom call-home events file. There are two types of .xml event files:

- ◆ EMC-generic events are modifications recommended by EMC.  
EMC provides an .xml file containing commonly requested modifications to the default call-home events.
- ◆ Customer-specific events are events modified to meet a specific customer requirement.  
EMC provides a custom events file developed by EMC engineering and applied by EMC Technical Support.

If no file is specified, this command removes both custom call-home events files.

The specified file is not deleted from the management server. When a custom events file is removed, the default events file LIC.xml is applied.

Use the **notifications call-home import-event-modifications** command to re-import the file.

**Example** In the following example, the specified customer-specific call-home events file is removed:

```
Vplexcli:/notifications/call-home> remove-event-modifications  
--customer-specific
```

```
The Customer-Specific call-home event modifications will be removed.  
Do you want to proceed? (Yes/No) y
```

**Example** In the following example, The EMC-generic call-home events file is removed:

```
Vplexcli:/notifications/call-home> remove-event-modifications  
--emc-generic
```

```
The EMC-generic call-home event modifications will be removed.  
Do you want to proceed? (Yes/No) yes
```

**Example** In the following example, both call-home events files are removed:

```
Vplexcli:/notifications/call-home> remove-event-modifications
```

```
The customer-specific and EMC-generic call-home event modifications  
will be removed.
```

```
Do you want to proceed? (Yes/No) yes
```

- See also**
- ◆ [notifications call-home import-event-modifications on page 359](#)
  - ◆ [notifications call-home view-event-modifications on page 363](#)
  - ◆ [notifications call-home test on page 365](#)

## notifications call-home view-event-modifications

Displays any customized call-home events.

**Contexts** All contexts.

In /notifications context, command is **call-home-view-event-modifications**

In /notifications/call-home context, command is **view-event-modifications**.

**Syntax** notifications call-home view-event-modifications  
[-c|--customer-specific]  
[-e|--emc-generic]

**Arguments** **Optional arguments**  
[-c|--customer-specific] - Displays customer specific modifications.  
[-e|--emc-generic] - Displays EMC generic modifications.

**Description** If event modifications are applied to call-home events, this command displays those events whose call-home events have been modified.

If the same event is modified by both the customer-specific and the EMC-generic events files, the setting in the customer-specific file overrides the entry in the EMC-generic file.

Use this command with no arguments to display a summary of all event modifications.

Use this command with the -c or -e arguments to display a summary of only the customer-specific or EMC generic modified events.

Use the --verbose argument to display detailed information.

**Example** Display a summary of event modifications:

```
Vplexcli:/notifications/call-home> view-event-modifications
```

```
EMC-generic events :
```

```
-----  
event code  name                severity  
-----  
0x8a023001  amf_1_WARNING                 INFO
```

```
Customer-specific events :
```

```
-----  
event code  name                                severity  
-----  
0x8a530006  CWS_EVENT_HEALTHY_CONNECTIVITY_INTERVAL_CHANGED  WARNING  
0x8a530005  CWS_EVENT_SILENCE_THRESHOLD_CHANGED              INFO
```

**Example** In the following example, the same event is modified by both the customer-specific and the EMC-generic events files. The setting in the customer-specific file overrides the setting in the EMC-generic file:

```
Vplexcli:/notifications/call-home> view-event-modifications
```

```
EMC-generic events :
```

```
-----  
event code  name                severity  
-----  
0x8a2d6025  SCSI_IT_LOST        WARNING  
0x8a2c901c  scom_28_CRIT        ERROR
```

Customer-specific events :

```
-----  
event code  name          severity  
-----  
0x8a2d6025  SCSI_IT_LOST  WARNING  
0x8a029060  amf_96_CRIT  CRITICAL
```

**Example** Use the **--verbose** argument to display detailed information about customer-specific event modifications:

```
VPlexcli:/notifications/call-home> view-event-modifications --customer-specific --verbose
```

Customer-specific events :

```
-----
```

For Event Code 0x8a530006:

Modified Fields : [Severity, Resend Timeout, Default Event, Customer Description]

```
name:                CWS_EVENT_HEALTHY_CONNECTIVITY_INTERVAL_CHANGED  
severity:            WARNING  
component:           cws  
locale:              CLUSTER  
obsolete event:      NO  
resend timeout:      60  
default event:       TRUE  
threshold count:     0  
threshold interval:  0  
customer description: CUSTOM Cluster Witness Server Healthy Connectivity Interval is changed.  
format string:       Cluster Witness Server Healthy Connectivity Interval is changed from  
%u to %u seconds
```

For Event Code 0x8a530005:

```
.  
. .
```

- See also**
- ◆ [notifications call-home remove-event-modifications on page 361](#)
  - ◆ [notifications call-home import-event-modifications on page 359](#)

---

## notifications call-home test

Sends a test event through call-home.

**Contexts** All contexts.

In /notifications context, command is **call-home test**.

In /notifications/call-home context, command is **test**.

**Syntax** `notifications call-home test`

**Description** Call-home can be configured to send events to EMC Support and/or one or more recipients in your organization.

Use this command to send a test event to the configured recipients. VPLEX sends the test call-home within 1 minute of running this command.

If call-home is configured to send event notifications to personnel in your organization, check the e-mail account(s) specified to receive notifications to verify the test event arrived.

If call-home is configured to send event notifications to EMC, contact EMC Support to verify that the test event arrived.

Use the **set** command to enable/disable call-home notifications. Refer to [Enable/disable call-home notifications on page 445](#).

Use the **ls notifications/call-home** command to verify that call-home is enabled:

```
Vplexcli: /> ls /notifications/call-home

/notifications/call-home:

Attributes:
Name      Value
-----  -----
enabled   true
```

**Example** Use the **call-home test** command to send a test call-home event to the configured recipients:

```
Vplexcli:/notifications/call-home> call-home test
call-home test was successful.
```

**See also**

- ◆ [configuration event-notices-reports config on page 121](#)
- ◆ [configuration event-notices-reports reset on page 124](#)
- ◆ [notifications snmp-trap create on page 366](#)
- ◆ [set on page 444](#)

## notifications snmp-trap create

Creates an SNMP trap sink for call-home events.

**Contexts** All contexts.

In /notifications context, command is **snmp-trap create**.

In /notifications/call-home/snmp-traps context, command is **create**.

**Syntax** notifications snmp-trap create  
[-n|--name] trap name

**Arguments** **Required arguments**  
[-n|--name] trap name - Name of the SNMP trap sink.

**Description** The SNMP trap does not start automatically.

To start the SNMP trap, do the following:

- ◆ Use the **set** command to set the IP address of the remote-host.
- ◆ Use the **set** command to set the **started** attribute to **true**.

### Example

```
VPlexcli:/notifications/call-home/snmp-traps> notifications snmp-trap create Test
VPlexcli:/notifications/call-home/snmp-traps> cd /Test
```

```
VPlexcli:/notifications/call-home/snmp-traps/Test> ll
```

| Name             | Value  |
|------------------|--------|
| community-string | public |
| remote-host      | -      |
| remote-port      | 162    |
| started          | false  |

```
VPlexcli:/notifications/call-home/snmp-traps/Test> set remote-host 10.6.213.39
VPlexcli:/notifications/call-home/snmp-traps/Test> set started true
```

- See also**
- ◆ [notifications call-home test on page 365](#)
  - ◆ [notifications snmp-trap destroy on page 367](#)
  - ◆ [set on page 444](#)

---

## notifications snmp-trap destroy

Destroys one or more SNMP traps.

**Contexts** All contexts.

In /notifications context, command is **snmp-trap destroy**.

In /notifications/call-home/snmp-traps context, command is **destroy**.

**Syntax** `notifications snmp-trap destroy`  
`[-s|--snmp-trap] trap name`  
`[-f|--force]`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-s|--snmp-trap] trap name` - Name of the SNMP trap sink to destroy.

`[-f|--force]` - Destroy an SNMP trap sink that has been started.

**Description** The `--force` argument is required to destroy an SNMP trap sink that has been started.

**Example** `Vplexcli:/notifications/call-home/snmp-traps> notifications snmp-trap destroy Test`

**See also** ♦ [notifications snmp-trap create on page 366](#)

---

## plugin addurl

Adds an URL to the plug-in search path.

**Contexts** All contexts.

**Syntax** `plugin addurl  
[-u|--urls] url,url...`

**Arguments** **Required arguments**  
`[-u|--urls] url,url...` - A list of URLs to add to the search path. Entries must be separated by commas.

**Description** Note: The plugin commands are not intended for customer use.

Plug-ins extend the class path of the CLI. Plug-ins support dynamic addition of functionality. The plugin search path is used by the “[plugin register](#)” command.

**See also**

- ◆ [plugin listurl on page 369](#)
- ◆ [plugin register on page 370](#)

---

## plugin listurl

Lists URLs currently in the plugin search path.

**Contexts** All contexts.

**Syntax** `plugin listurl`

**Description** The search path URLs are those locations added to the plugin search path using the “[plugin addurl](#)” command.

---

**Note:** The plugin commands are not intended for customer use.

---

**Example** `VPlexcli:/> plugin listurl`

```
file:/opt/emc/VPlex/jython2.2/LibExt/AutoBundles/prodscripts.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/bin/commons-daemon.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/bin/bootstrap.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/bin/tomcat-juli.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-i18n-es.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-juli-adapters.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/catalina-tribes.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/servlet-api.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-coyote.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/realms-adapter.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/catalina-ha.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/jasper-jdt.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/catalina.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/catalina-ant.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/jsp-api.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/annotations-api.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/jasper-el.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/jasper.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-i18n-ja.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/el-api.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-i18n-fr.jar,  
file:/opt/emc/VPlex/apache-tomcat-6.0.x/lib/tomcat-dbcp.jar
```

**See also**

- ◆ [plugin addurl on page 368](#)
- ◆ [plugin register on page 370](#)

---

## plugin register

Registers a shell plugin by class name.

**Contexts** All contexts.

**Syntax** `plugin register [-c|--classes] list of class names, list of class names ...`

**Arguments** **Required arguments**  
`[-c|--classes] list of class names, list of class names...` - A list of plugin classes. Entries must be separated by commas.

**Description** Plugin class is found in the default classpath, or in locations added using the “[plugin addurl](#)” command.

Plug-ins add a batch of commands to the CLI, generally implemented as a set of one or more Jython modules.

---

**Note:** The plugin commands are not intended for customer use.

---

**See also**

- ◆ [plugin addurl on page 368](#)
- ◆ [plugin listurl on page 369](#)

---

## popd

Pops the top context off the stack, and changes the current context to that context.

**Contexts** All contexts.

**Syntax** `popd`

**Description** If the context stack is currently empty, an error message is displayed.

**Example** In the following example:

- ◆ The **pushd** command adds a third context to the context stack. The output of the command displays the three contexts in the stack.
- ◆ The **popd** command removes the top (last added) context, changes the context to the next one in the stack, and the output displays the two remaining contexts:

```
Vplexcli:/engines/engine-1-1/directors/Cluster_1_Dir1A> pushd
/engines/engine-1-1/directors/Cluster_1_Dir1B
[/engines/engine-1-1/directors/Cluster_1_Dir1B,
 /engines/engine-1-1/directors/Cluster_1_Dir1A,
 /clusters/cluster-1/storage-elements/storage-arrays, /, /]

Vplexcli:/engines/engine-1-1/directors/Cluster_1_Dir1B> popd
[/engines/engine-1-1/directors/Cluster_1_Dir1B,
 /clusters/cluster-1/storage-elements/storage-arrays, /, /]

Vplexcli:/engines/engine-1-1/directors/Cluster_1_Dir1A>
```

**See also** ◆ [pushd on page 380](#)

## ptov describe-be-zoning

Lists the zones that should be created for physical to virtual insertion.

**Contexts** All contexts.

**Syntax**  
`ptov describe-be-zoning`  
`[-v|--virtualizedStorage] filename`  
`[-s|--zoningSummary] filename`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-v|--virtualizedStorage] filename` - Virtual storage file created by the “[ptov suggest-be-zoning-and-masking](#)” command that suggests back-end zoning and masking.

`[-s|--zoningSummary] filename` - Location where zoning summary file is output.

**Description** Uses the virtual storage file created by the “[ptov suggest-be-zoning-and-masking](#)” command to create a list of port pairs that should be zoned together.

The output is displayed in a format that is easy to read, and thus reduces the likelihood that network administrators make incorrect zoning decisions.

This information is also in the virtual storage file.

See [Using ptov commands on page 374](#) for an overview and the order in which the ptov commands should be used.

### Example

```
VPLEXcli: /> ptov describe-be-zoning --virtualizedStorage virtualized.txt
```

- See also**
- ◆ [ptov lun-mask-be-storage on page 373](#)
  - ◆ [ptov query-be-storage on page 374](#)
  - ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
  - ◆ [ptov suggest-ports-to-use on page 377](#)
  - ◆ [ptov verify-be-storage on page 378](#)
  - ◆ [ptov verify-be-zoning on page 379](#)

## ptov lun-mask-be-storage

Perform LUN masking for physical to virtual insertion.

**Contexts** All contexts.

**Syntax**

```
ptov lun-mask-be-storage
  [--smisport] port
  [--cimnamespace] cim namespace
  [-u|--user] username
  [-p|--password] password
  [-v|--virtualizedStorage] filename
  [-smishost] smishost
```

**Arguments** **Required arguments**

**[-u|--user] *username*** - Username to use to connect to the host.

**--smishost *smishost*** - \* SMI-S provider host to which to connect.

**Optional arguments**

**[--smistport] *port*** - SMI-S port for the host.

**[--cimnamespace] *cim namespace*** - CIM namespace in which to search.

**[-p|--password] *password*** - Password to use to connect to the host. If no password is provided, a prompt is displayed.

**[-v|--virtualizedStorage]** - File containing the storage information that will be added to the VPLEX by the virtualization.

**[-smishost] *smishost*** - \* SMI-S provider host to which to connect.

\* - argument is positional.

**Description** Uses the SMI-S host to perform the LUN masking described in the virtual storage file created by the “[ptov suggest-be-zoning-and-masking](#)”.



**CAUTION**

**If the LUN masking fails, all masking must be undone, including any portion that was successfully masked.**

See [Using ptov commands on page 374](#) for an overview and the order in which the ptov commands should be used.

**Example**

```
Vplexcli: />ptov lun-mask-be-storage 128.0.0.256 -u john -v virtualized.txt
```

- See also**
- ◆ [ptov describe-be-zoning on page 372](#)
  - ◆ [ptov query-be-storage on page 374](#)
  - ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
  - ◆ [ptov suggest-ports-to-use on page 377](#)
  - ◆ [ptov verify-be-storage on page 378](#)
  - ◆ [ptov verify-be-zoning on page 379](#)

## ptov query-be-storage

Query the existing storage setup and send the output to the specified file.

**Contexts** All contexts.

**Syntax**

```
ptov query-be-storage
  [--smisport] port
  [--cimnamespace] cim namespace
  [-u|--user] username
  [-p|--password] password
  [-e|--existingStorage] filename
  [-smishost] smishost
```

**Arguments** **Required arguments**

**[-u|--user] *username*** - Username to use to connect to the host.

**--smishost *smishost*** - \* SMI-S provider host to which to connect.

**Optional arguments**

**[--smistport] *port*** - SMI-S port for the host.

**[--cimnamespace] *cim namespace*** - CIM namespace in which to search.

**[-p|--password] *password*** - Password to use to connect to the host. If no password is provided, a prompt is displayed.

**[-e|--existingStorage]** - File where the configuration of already existing storage will be stored.

**[--smishost] *smishost*** - \* SMI-S provider host to which to connect.

\* - argument is positional.

**Description** Queries the storage and creates a file that describes the existing storage environment.

Output is an “existing storage file” used by other steps in the physical-to-virtual insertion process.

**Using ptoV commands**

VPLEX resides between storage and application hosts. The ptoV commands help plan and connect a VPLEX to existing storage. The order in which the commands are intended to be used is as follows:

- ◆ **“ptov query-be-storage”** - queries the storage and creates a file that describes the existing storage environment.

Output of the command consists of an existing storage file.

Users edit the file, removing any storage that should not be encapsulated as part of the initiator group being virtualized.

For example, if storage for application 1 is being virtualized, then storage associated with application 2 should be removed from the file.

- ◆ **“ptov suggest-ports-to-use”** - displays information about port load and suggests the least loaded ports.
- ◆ **“ptov suggest-be-zoning-and-masking”** - uses the edited existing storage file, and a user provided list of VPLEX back-end ports to suggest back end zoning and masking.

If less than 2 ports are specified, the command fails (it is not best practice to wire storage without redundant connections).

Output of the command is a virtual storage file.

- ◆ **“ptov describe-be-zoning”** - uses the virtual storage file created by the **“ptov suggest-be-zoning-and-masking”** command to create a list of port pairs that should be zoned together.

The output is displayed in a format that is easy to read, and thus reduces the likelihood that network administrators make incorrect zoning decisions.

This information is also in the virtual storage file.

- ◆ Plug in Fibre channel cables and configure switches in order to zone the network according to the information created by the previous steps.
- ◆ **“ptov verify-be-zoning”** - checks the zoning, and returns errors there are problems.

Output when errors are discovered includes which zones are expected but missing.

After correcting the errors, this command is re-run. This step is repeated until there are no errors.

- ◆ **“ptov lun-mask-be-storage”**- uses smi-s to perform the LUN masking described in the virtual storage file.
- ◆ If the LUN masking fails, all masking must be undone, including the any part that worked.
- ◆ **“ptov verify-be-storage”** - Verifies that the storage seen from the directors is consistent with the description in the virtual storage file.

### Example

```
Vplexcli:/> ptov query-be-storage 128.1.1.256 -u john --existingStorage to-virtualize.txt
```

- See also**
- ◆ [ptov describe-be-zoning on page 372](#)
  - ◆ [ptov lun-mask-be-storage on page 373](#)
  - ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
  - ◆ [ptov suggest-ports-to-use on page 377](#)
  - ◆ [ptov verify-be-storage on page 378](#)
  - ◆ [ptov verify-be-zoning on page 379](#)

## ptov suggest-be-zoning-and-masking

Given a description of a current physical storage configuration, output a file that suggests zoning and LUN masking.

**Contexts** All contexts.

**Syntax**

```
ptov suggest-be-zoning-and-masking
[-c|--cluster] cluster
[-t|--groupToVirtualize] filename
[-v|--virtualizedStorage] filename
[-p|--portsToUse] port,port...
```

**Arguments** **Required arguments**

**[-p|--portsToUse] port,port...** - List of back-end ports to use to encapsulate the storage. Entries must be separated by commas.

**Optional arguments**

**[-c|--cluster] cluster** - Cluster in which to insert.

**[-t|--groupToVirtualize] filename** - File containing only the storage information for storage groups that should be virtualized.

**[-v|--virtualizedStorage]** - File containing the storage information that will be added to the VPLEX by the virtualization.

**Description**

Uses the edited existing storage file created by “[ptov query-be-storage](#)”, and a user provided list of VPLEX back-end ports to suggest back end zoning and masking.

If less than two ports are specified, the command fails (it is not best practice to wire storage without redundant connections).

Output of the command is a virtual storage file.

See [Using ptov commands on page 374](#) for an overview and the order in which the ptov commands should be used.

**Example**

```
VPLexcli:/> ptov suggest-be-zoning-and-masking --portsToUse A2-FC00, A3-FC00 -c cluster-1
--groupToVirtualize to-virtualize.txt --virtualizedStorage virtualized.txt
```

**See also**

- ◆ [ptov describe-be-zoning on page 372](#)
- ◆ [ptov lun-mask-be-storage on page 373](#)
- ◆ [ptov query-be-storage on page 374](#)
- ◆ [ptov suggest-ports-to-use on page 377](#)
- ◆ [ptov verify-be-storage on page 378](#)
- ◆ [ptov verify-be-zoning on page 379](#)

---

## ptov suggest-ports-to-use

Suggest which ports would be best to use for the virtualization.

**Contexts** All contexts.

**Syntax**  
`ptov suggest-ports-to-use  
[-c|--cluster] cluster  
[--planned] plan,plan...`

**Arguments** **Required arguments**  
None.

**Optional arguments**

`[-c|--cluster] cluster` - Cluster in which to insert.

`[--planned] plan,plan...` - Storage files for initiator groups that have been planned but not yet zoned and masked.

**Description** Displays information about port load and suggests the least loaded ports.

See [Using ptov commands on page 374](#) for an overview and the order in which the ptov commands should be used.

**Example** `Vplexcli: /> ptov suggest-ports-to-use -c cluster-1`

- See also**
- ◆ [ptov describe-be-zoning on page 372](#)
  - ◆ [ptov lun-mask-be-storage on page 373](#)
  - ◆ [ptov query-be-storage on page 374](#)
  - ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
  - ◆ [ptov verify-be-storage on page 378](#)
  - ◆ [ptov verify-be-zoning on page 379](#)

## ptov verify-be-storage

Verify connectivity for the back-end stage of the physical to virtual insertion

**Contexts** All contexts.

**Syntax**

```
ptov verify-be-storage
  [--smisport] port
  [--cimnamespace] cim namespace
  [-u|--user] username
  [-p|--password] password
  [-v|--virtualizedStorage] filename
  [-smishost] smishost
```

**Arguments** **Required arguments**

**[-u|--user] *username*** - Username to use to connect to the host.

**--smishost *smishost*** - \* SMI-S provider host to which to connect.

**Optional arguments**

**[--smistport] *port*** - SMI-S port for the host.

**[--cimnamespace] *cim namespace*** - CIM namespace in which to search.

**[-p|--password] *password*** - Password to use to connect to the host. If no password is provided, a prompt is displayed.

**[-v|--virtualizedStorage]** - File containing the storage information that will be added to the VPLEX by the virtualization.

**[-c|--cluster] *cluster*** - Cluster in which to insert.

**[--smishost] *smishost*** - \* SMI-S provider host to which to connect.

\* - argument is positional.

**Description**

Verifies that the storage seen from the directors is consistent with the description in the virtual storage file created by “[ptov suggest-be-zoning-and-masking](#)”.

See [Using ptoV commands on page 374](#) for an overview and the order in which the ptoV commands should be used.

**Example**

```
Vplexcli: /> ptoV verify-be-storage 128.0.0.256 -u John -c cluster-1 -v virtualized.txt
```

**See also**

- ◆ [ptov describe-be-zoning on page 372](#)
- ◆ [ptov lun-mask-be-storage on page 373](#)
- ◆ [ptov query-be-storage on page 374](#)
- ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
- ◆ [ptov suggest-ports-to-use on page 377](#)
- ◆ [ptov verify-be-zoning on page 379](#)

---

## ptov verify-be-zoning

Verify connectivity for the back-end stage of the physical to virtual insertion.

**Contexts** All contexts.

**Syntax**  
`ptov verify-be-zoning  
[-v|--virtualizedStorage] filename  
[-c|--cluster] cluster`

**Arguments** **Required arguments**  
None.

**Optional arguments**

`[-v|--virtualizedStorage]` - File containing the storage information that will be added to the VPLEX by the virtualization.

`[-c|--cluster] cluster` - Cluster in which to insert.

**Description** Checks the zoning, and returns errors if there are problems.

Error messages include which zones are expected but missing.

After correcting the errors, re-run the **ptov verify-be-zoning** command. Repeat until there are no errors.

When there are no errors, proceed to “[ptov lun-mask-be-storage](#)”.

See [Using ptov commands on page 374](#) for an overview and the order in which the ptov commands should be used.

**Example**

```
Vplexcli:/> ptov verify-be-zoning -c cluster-1 --virtualizedStorage virtualized.txt
```

- See also**
- ◆ [ptov describe-be-zoning on page 372](#)
  - ◆ [ptov lun-mask-be-storage on page 373](#)
  - ◆ [ptov query-be-storage on page 374](#)
  - ◆ [ptov suggest-be-zoning-and-masking on page 376](#)
  - ◆ [ptov suggest-ports-to-use on page 377](#)
  - ◆ [ptov verify-be-storage on page 378](#)

## pushd

Pushes the current context onto the context stack, and then changes the current context to the given context.

**Contexts** All contexts.

**Syntax** `pushd`  
`[-c|--context] context`

**Arguments** **Required arguments**  
None.

**Optional arguments**  
`[-c|--context] context` - The context to push onto the context stack.

**Description** Adds the context to the context stack.

If no context is supplied, and there is a context on the stack, the current context is exchanged with the top-of-stack context.

Use the “[popd](#)” command to remove the topmost context from the context stack.

**Example** Starting in the root context:

```
VPlexcli: />
```

Use the **pushd** command to push the first context onto the context stack:

```
VPlexcli: /> pushd  
/clusters/cluster-1/storage-elements/storage-arrays/  
[/clusters/cluster-1/storage-elements/storage-arrays, /, /]
```

Use the **pushd** command to push a second context onto the context stack:

```
VPlexcli: /clusters/cluster-1/storage-elements/storage-arrays> pushd  
/engines/engine-1-1/directors/Cluster_1_Dir1A/  
[/engines/engine-1-1/directors/Cluster_1_Dir1A,  
/clusters/cluster-1/storage-elements/storage-arrays, /, /]
```

Now, there are two contexts on the context stack. Use the **pushd** command to toggle between the two contexts:

```
VPlexcli: /engines/engine-1-1/directors/Cluster_1_Dir1A> pushd  
[/clusters/cluster-1/storage-elements/storage-arrays,  
/engines/engine-1-1/directors/Cluster_1_Dir1A, /, /]
```

```
VPlexcli: /clusters/cluster-1/storage-elements/storage-arrays> pushd  
[/engines/engine-1-1/directors/Cluster_1_Dir1A,  
/clusters/cluster-1/storage-elements/storage-arrays, /, /]
```

```
VPlexcli: /engines/engine-1-1/directors/Cluster_1_Dir1A>
```

**See also** ♦ [popd on page 371](#)

## rebuild set-transfer-size

Changes the transfer-size of the given devices.

**Contexts** All contexts.

**Syntax** `rebuild set-transfer-size  
[-r|--devices] context path,context path...  
[-l|--limit] limit`

**Arguments** **Required arguments**

`[-r|--devices] context path,context path...` - \* List of one or more devices for which to change the transfer size. Wildcards are permitted. Entries must be separated by commas.

`[-l|--limit] limit` - \* Transfer size in bytes. Maximum number of bytes to transfer as one operation per device. Specifies the size of read sector designated for transfer in cache. Setting this value smaller implies more host I/O outside the transfer boundaries. Setting the value larger may result in faster transfers. Valid values must be multiples of 4K.

Range: 40K-128M.

See [About transfer-size on page 53](#).

\* - argument is positional.

**Description** If the target device(s) are rebuilding when this command is issued, the rebuild is paused and resumed using the new transfer-size.

**Note:** If there are queued rebuilds, the rebuild may not resume immediately.

**Example** Set the transfer-size on a specified device to 1M:

```
Vplexcli:/> rebuild set-transfer-size --devices /clusters/cluster-1/devices/testdevice --limit 1M
```

Set the transfer-size for all devices to 2M:

```
Vplexcli:/> rebuild set-transfer-size /clusters/*/devices/* 2M
```

Set the transfer-size for all distributed devices to 10K:

```
Vplexcli:/distributed-storage/distributed-devices> rebuild set-transfer-size * 10k
```

**See also**

- ◆ [rebuild show-transfer-size on page 382](#)
- ◆ [rebuild status on page 383](#)

## rebuild show-transfer-size

Shows the transfer-size of specified RAID 1 devices.

**Contexts** All contexts.

**Syntax** `rebuild show-transfer-size  
[-r|--devices] context path`

### Optional arguments

`[-r|--devices] context path...` - List of one or more RAID 1 devices for which to display the transfer size. Entries must be separated by commas. Wildcards are permitted.

**Example** Display the rebuild transfer size for a specified device:

```
Vplexcli:/> rebuild show-transfer-size TestDevice  
device name      transfer-size  
-----  
TestDevice       2M
```

Display rebuild transfer size for selected devices:

```
Vplexcli:/> rebuild show-transfer-size dd_0*  
device name      transfer-size  
-----  
dd_00           2M  
dd_01           2M  
dd_02           2M  
.  
.  
.
```

Display rebuild transfer size for all distributed devices:

```
Vplexcli:/> rebuild show-transfer-size *  
device name      transfer-size  
-----  
TestDevice       2M  
dd_00            2M  
dd_01            2M  
dd_02            2M  
.  
.  
.
```

**See also**

- ◆ [rebuild set-transfer-size on page 381](#)
- ◆ [rebuild status on page 383](#)

## rebuild status

Displays all global and cluster-local rebuilds along with their completion status.

**Contexts** All contexts.

**Syntax** `rebuild status`  
`--show-storage-volumes`

### Optional arguments

**--show-storage-volumes** - Displays all storage volumes that need to be rebuilt, both active and queued. If not present, only the active rebuilds are displayed.

**Description** Completion status is listed as:

rebuilt/total (complete%)

**Example** Check rebuild status from storage volume context:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> rebuild status  
[1] disks marked for rebuild
```

```
Global rebuilds:  
No active global rebuilds.
```

```
cluster-1 local rebuilds:  
device      rebuild type  rebuilder director  rebuilt/total  percent finished  throughput  ETA  
-----  
test3313_r1  full         s10_428f          1.23G/4G      30.81%         90.1M/s
```

Check rebuild status from distributed-device-components/*volume*/components/*context*, and display storage volumes that need to be rebuilt:

```
Vplexcli:/distributed-storage/distributed-devices/testvol1/distributed-device-components/C2testvol10000/components> rebuild status --show-storage-volumes
```

```
StorageVolumes marked for rebuild:  
cluster-2:  
  extent_60060160639028006413c641e2a7e011_1
```

```
[1] storage_volumes marked for rebuild
```

```
Global rebuilds:  
device      rebuild type  rebuilder director  rebuilt/total  percent finished  throughput  ETA  
-----  
testvol1    full         s1_220d_spa      4.06G/11.2G   36.17%         9.94M/s     12.3min
```

```
Local rebuilds:  
No active local rebuilds.
```

**See also** ♦ [rebuild show-transfer-size on page 382](#)

## remote-clusters add-addresses

Adds one or more address:subnet mask configurations for the specified remote-cluster entry for this cluster.

**Contexts** clusters/*cluster*/cluster-connectivity

**Syntax** remote-clusters add-addresses  
[-c|--cluster] *cluster*  
[-d|--default]  
[-r|--remote-cluster] *remote-cluster*  
[-a|--addresses] *IP address:subnet mask, IP address:subnet mask...*

**Arguments** [-c|--cluster] *cluster* - \* Context path for the cluster whose connectivity configuration is to be modified. Typically, the cluster directly above the current context.

[-d|--default] - Applies the default configuration of remote cluster and addresses. Default values are determined by the cluster-address attribute of the active subnets from all remote clusters. If this argument is used, no values are required for --remote-cluster or --addresses.

[-r|--remote-cluster] *remote-cluster* - Context path for the remote-cluster configuration entry to modify. Can not be the same context specified by the --cluster argument. This argument is not required when --default is used. If this argument is used, the --addresses argument is also required.

[-a|--addresses] *IP address:subnet mask, IP address:subnet mask...* - A list of remote IP address:subnet mask pairs to add for the specified --remote-cluster. This argument is not required when --default is used. If this argument is used, the --remote-cluster argument is also required.

\* argument is positional.

**Description** Creates a list of reachable IP addresses for a remote cluster.

**Example** Create a default list of reachable IP addresses (using the cluster-address attribute of the active subnets of remote clusters) for all remote clusters:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity> remote-clusters add-addresses --default
```

**Example** Create a list of specific reachable IP addresses for a remote cluster:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity> remote-clusters add-addresses  
--remote-cluster cluster-2 --addresses 42.29.20.254:1100, 63.12.40.254:11000
```

**See also** ♦ [remote-clusters clear-addresses on page 385](#)

## remote-clusters clear-addresses

Clear one, several or all address:subnet mask pairs for the specified remote-cluster entry for this cluster.

**Contexts** clusters/*cluster*/cluster-connectivity

**Syntax** remote-clusters clear-addresses  
[-c|--cluster] *cluster*  
[-r|--remote-cluster] *remote-cluster*  
[-a|--addresses] IP address:subnet mask, IP address:subnet mask...

**Arguments** [-c|--cluster] *cluster* - \* Context path for the cluster whose connectivity configuration is to be modified. Typically, the cluster directly above the current context.

[-r|--remote-cluster] *remote-cluster* - \* Context path for the remote-cluster configuration entry to modify. Can not be the same context specified by the --cluster argument.

[-a|--addresses] *IP address:subnet mask, IP address:subnet mask...* - A list of one or more remote ip-address:subnet mask entries to remove for the specified --remote-cluster. If no ip-address:subnet mask entry is specified, all entries are removed for the specified --remote-cluster.

\* argument is positional.

**Description** Clears one or more reachable addresses at a remote cluster.

**Example** Clear a specific address:subnet mask pair for the specified remote cluster:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity> remote-clusters clear-addresses  
--remote-cluster cluster-2 --addresses 42.29.20.254:1100, 63.12.40.254:11000
```

**Example** Clear all reachable addresses for the remote cluster:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity> remote-clusters clear-addresses  
--remote-cluster cluster-2
```

**See also** ♦ [remote-clusters add-addresses on page 384](#)

## report aggregate-monitors

Aggregate the reports generated by the **report create-monitors** or **monitor** commands.

**Contexts** All contexts.

**Syntax** `report aggregate-monitors  
[-d|--directory] directory`

### Optional arguments

`[-d|--directory] directory`- Directory in which to create the .csv files.

Default directory path: /var/log/VPLEX/cli/reports/ on the management server.

**Description** The reports are aggregated by cluster.

An aggregate report is generated for:

- ◆ Each cluster
- ◆ Each type of per director report in the specified directory

This command assumes that the per director reports have filenames with the following format:

```
<report type>ReportMonitor_<director>.csv
```

All other files in the directory will be ignored.

Aggregate report filenames are in the following format:

```
<report type>Performance_<cluster>.csv
```

If an aggregate report already exists, it will be overwritten.

### Example

```
Vplexcli:/> report aggregate-monitors  
Aggregating port reports at cluster cluster-2 ...  
Aggregating port reports at cluster cluster-1 ...  
Aggregating volume reports at cluster cluster-2 ...  
Aggregating volume reports at cluster cluster-1 ...  
.  
.  
.
```

Output files saved in /var/log/VPLEX/cli/reports/:

- ◆ diskPerformance\_cluster-1.csv
- ◆ diskPerformance\_cluster-2.csv
- ◆ portPerformance\_cluster-1.csv
- ◆ portPerformance\_cluster-2.csv
- ◆ volumePerformance\_cluster-1.csv
- ◆ volumePerformance\_cluster-2.csv

### See also

- ◆ [report capacity-arrays on page 387](#)
- ◆ [report capacity-clusters on page 390](#)
- ◆ [report capacity-hosts on page 391](#)
- ◆ [report create-monitors on page 392](#)
- ◆ [report poll-monitors on page 395](#)

## report capacity-arrays

Generates a capacity report.

**Contexts** All contexts.

**Syntax**  
report capacity-arrays  
[-t|--tier-regex] *regular expression*  
[-d|--directory] *directory*

### Optional arguments

**[-t | -tier-regex] *regular expression*** - Regular expression which when applied to the storage-volume name, returns the tier ID in a group. Most expressions must be enclosed in quotes.

Default: value of /system-defaults::tier-regular-expression

**[-d | -directory] *directory*** - Directory in which to create the csv files. Output is written to two files:

- ◆ File for local storage: CapacityArraysLocal.csv
- ◆ File for shared storage: CapacityArraysShared.csv.

Default directory path: /var/log/VPlex/cli/reports/\* on the management server.

**Description** Generates a capacity report for all the storage in a VPlex, grouped by storage arrays.

This command assumes the following:

- ◆ All storage volumes in a storage array have the same tier value.
- ◆ The tier is indicated in the storage-volume name. The tier attribute in the virtual volumes context is ignored.

If a file is specified, output is formatted as:

```
<time>,<cluster name>,<array name>,<tier string>,<alloc>,<unalloc devices>,<unalloc storage-volumes> <time>,<alloc>,<unalloc devices>
```

If the file(s) already exist, the report is appended to the end of the file(s).

**Note:** Tier IDs are required to determine the tier of a storage-volume/storage array. Storage volumes that do not contain any of the specified IDs are given the tier value 'no-tier'.

The report is separated into two parts: local storage and shared storage.

- ◆ **Local storage** is accessible only from the same cluster where the storage is physically located. Information in the report for local storage includes:
  - Cluster id
  - Storage array
  - Tier - the tier of the storage array
  - Allocated - storage that is visible through a view (exported)
  - Unallocated-device - storage that is in devices, but not visible from a view. For example, a virtual volume that has not been exported or free space in a device that is not part of a virtual volume.
  - Unallocated-storage-volume - storage in unused storage volumes.
- ◆ **Shared storage** is accessible from clusters other than where it is physically located (distributed and remote virtual volumes). Information in the report for shared storage includes:

- allocated - storage that is visible through a view (exported)
- unallocated-device - storage that is in devices, but not visible from a view. For example, a virtual volume that has not been exported or free space in a device that is not part of a virtual volume.

---

**Note:** Unreachable storage volumes are not included in 'unallocated-storage-volume'. If the storage array of a virtual volume/device cannot be determined, its capacity is allocated to the 'unknown array'.

---

There is no 'tier' indicator for shared storage because the tiers may be different for each mirror of a distributed-device.

There are no shared storage volumes. Storage Volumes are only locally accessible and are part of the cluster allocation.

**Display the report file** To display the raw report file, do the following:

1. Exit to the management server:

```
Vplexcli:/> exit
Connection closed by foreign host.
service@ManagementServer:~>
```

2. Navigate to the VPLEX CLI reports directory (or the specified output directory):

```
service@ManagementServer:~> cd /var/log/Vplex/cli/reports
service@ManagementServer:/var/log/Vplex/cli/reports> ll
total 48
-rw-r--r-- 1 service users 2253 2010-08-12 15:46
  CapacityArraysLocal.csv
-rw-r--r-- 1 service users 169 2010-08-12 15:46
  CapacityArraysShared.csv
.
.
.
```

3. Use the `cat filename` command to display the file:

```
service@ManagementServer:/var/log/Vplex/cli/reports> cat CapacityArraysLocal.csv
Time, Cluster name, Array name, Tier string, Allocated volumes (GiB), Unalloc devices (GiB),
Unalloc storage_volumes (GiB)
2010-06-21 16:00:32, cluster-1, EMC-0x00000000192601378, no-tier, 0, 0, 5666242560000
2010-06-21 16:00:32, cluster-1, EMC-0x00000000192601852, no-tier, 0, 0, 5292530073600
```

### Example

```
Vplexcli:/> report capacity-arrays
Local Storage (in GiB):
  cluster-1:
    EMC-0x00000000192601378:
      tier:                               no-tier
      alloc:                               0.000
      unalloc devices:                     0.000
      unalloc storage-volumes:             5277.100
    EMC-0x00000000192601852:
      tier:                               no-tier
      alloc:                               0.000
      unalloc devices:                     0.000
      unalloc storage-volumes:             4929.053
.
.
.
Shared Storage (in GiB):
  alloc:                                  2192.468
  unalloc devices:                        0.000
```

Tier summary (local, by cluster):

| cluster-1 |                 | Unallocated   |                      | Unallocated   |                      |
|-----------|-----------------|---------------|----------------------|---------------|----------------------|
| Tier      | Allocated (GiB) | Devices (GiB) | StorageVolumes (GiB) | Devices (GiB) | StorageVolumes (GiB) |
| no-tier   | 9.000           | 3.000         | 10596.152            |               |                      |

| cluster-2 |                 | Unallocated   |                      | Unallocated   |                      |
|-----------|-----------------|---------------|----------------------|---------------|----------------------|
| Tier      | Allocated (GiB) | Devices (GiB) | StorageVolumes (GiB) | Devices (GiB) | StorageVolumes (GiB) |
| no-tier   | 2.000           | 4.000         | 8467.749             |               |                      |

"report capacity-arrays" was not able to extract the tier id from the following storage-volumes/devices. Please ensure that the tier-regex contains 1 capture group.

tier-regex: ^[^\\_]+([HL])\_.\$

storage-volumes/devices:

```
Symm1723_1FC
CX4_lun0
base_volume
Symm1852_1C0
VPD83T3:60000970000192601378533030313530
Symm1852_5C0
VPD83T3:60000970000192601378533030313538
```

- See also**
- ◆ [report capacity-clusters on page 390](#)
  - ◆ [report capacity-hosts on page 391](#)

## report capacity-clusters

Generates a capacity report for every cluster.

**Contexts** All contexts.

**Syntax** `report capacity-clusters  
[-d|--directory] directory`

### Optional arguments

`[-d|--directory] directory`- Directory in which to create the csv files. Output is written to a file named CapacityClusters.csv.

Default directory path: /var/log/VPlex/cli/reports/ on the management server.

If the `--verbose` argument is used:

- ◆ Storage\_volumes and volumes are listed before the summary is printed.
- ◆ Two additional files are created: CapacityClustersVolumes.csv and CapacityClustersStorageVolumes.csv.

**Description** Information includes:

- ◆ Unclaimed storage-volume capacity in GB.
- ◆ Number of unclaimed storage volumes.
- ◆ Claimed storage-volume capacity in GB.
- ◆ Number of claimed storage volumes.
- ◆ Used storage-volume capacity in GB.
- ◆ Number of used storage volumes.
- ◆ Unexported virtual volume capacity in GB.
- ◆ Number of unexported virtual volumes.
- ◆ Exported virtual volume capacity in GB.
- ◆ Number of exported virtual volumes.

### Example

```
VPlexcli:/> report capacity-clusters  
Cluster, Unclaimed disk capacity (GiB), Unclaimed storage_volumes, Claimed disk capacity (GiB),  
Claimed storage_volumes, Used storage-volume capacity (GiB), Used storage_volumes, Unexported  
volume capacity (GiB), Unexported volumes, Exported volume capacity (GiB), Exported volumes  
cluster-1, 5705.13, 341, 7947.68, 492, 360.04, 15, 3.00, 3, 2201.47, 27  
cluster-2, 5337.10, 328, 7995.69, 495, 2478.45, 137, 20.00, 3, 2178.46, 25  
VPlexcli:/> report capacity-clusters --verbose  
Cluster, StorageVolume Name, VPD83 ID, Capacity, Use, Vendor  
cluster-1,CX4_Logging, VPD83T3:6006016021d02500e6d58bab2227df11, 80G, used, DGC  
cluster-1,CX4__M0, VPD83T3:6006016021d02500be83caff0427df11, 90G, -data, DGC  
cluster-1,CX4__M1, VPD83T3:6006016021d02500bf83caff0427df11, 90G, claimed, DGC  
cluster-1,CX4_lun0, VPD83T3:6006016021d0250026b925ff60b5de11, 10G, used, DGC  
.  
.  
.
```

**. See also** ◆ [report capacity-arrays on page 387](#)  
◆ [report capacity-hosts on page 391](#)

## report capacity-hosts

Generates a host capacity report.

**Contexts** All contexts.

**Syntax** `report capacity-hosts`  
`[-d|--directory] directory`

### Optional arguments

**[-d|--directory] *directory***- Directory in which to create the csv files. Output is written to a file named CapacityHosts.csv.

Default directory path: /var/log/VPlex/cli/reports/ on the management server.

If the **--verbose** argument is used, an additional file is created:  
CapacityHostsViews.csv

**Description** Information includes:

- ◆ Number of views.
- ◆ Total exported capacity in GB.
- ◆ Number of exported virtual volumes per cluster.

### Example

```
VPlexcli:/> report capacity-hosts
Cluster, Views, Exported capacity (GiB), Exported volumes
cluster-1, 2, 2209.47, 28
cluster-2, 1, 2178.46, 25
```

The **--verbose** argument prints view details:

```
VPlexcli:/> report capacity-hosts --verbose
Cluster, View name, Initiator ports, Target ports, Volumes, Capacity
cluster-1, LicoJ013, LicoJ013_hba3 LicoJ013_hba2 LicoJ013_hba1 LicoJ013_hba0,
P000000003CA00147-A1-FC00 P000000003CA00147-A1-FC02 P000000003CB00147-B0-FC02
P000000003CB00147-B1-FC00 P000000003CA00147-A0-FC00 P000000003CB00147-B0-FC00
P000000003CA00147-A0-FC02 P000000003CB00147-B1-FC02, dev_sym1723_1FC_vol, 8G
cluster-1, LicoJ009, LicoJ009_hba1 LicoJ009_hba0 LicoJ009_hba3 LicoJ009_hba2,
P000000003CA00147-A1-FC02 P000000003CB00147-B0-FC02 P000000003CA00147-A0-FC02
P000000003CB00147-B1-FC02, dd_09_vol dev_sym1723_1FC_vol TestDisDevice_vol dd_15_vol dd_16_vol
dd_20_vol dd_22_vol dd_18_vol dd_02_vol dd_12_vol dd_07_vol dd_19_vol dd_14_vol dd_13_vol
dd_04_vol dd_08_vol dd_11_vol dd_05_vol base0_vol dd_10_vol dd_23_vol dd_01_vol dd_00_vol
dd_17_vol dd_06_vol dd_03_vol dd_21_vol, 2.15T
cluster-2, LicoJ010, LicoJ010_hba1 LicoJ010_hba0 LicoJ010_hba3 LicoJ010_hba2,
P000000003CB000E6-B1-FC00 P000000003CB001CB-B0-FC00 P000000003CA001CB-A0-FC00
P000000003CA000E6-A0-FC00 P000000003CB001CB-B1-FC00 P000000003CB000E6-B0-FC00
P000000003CA001CB-A1-FC00 P000000003CA000E6-A1-FC00, base01_vol dd_09_vol dd_15_vol dd_16_vol
dd_20_vol dd_22_vol dd_18_vol dd_02_vol dd_12_vol dd_19_vol dd_07_vol dd_14_vol dd_13_vol
dd_04_vol dd_08_vol dd_11_vol dd_05_vol dd_10_vol dd_23_vol dd_01_vol dd_00_vol dd_17_vol
dd_06_vol dd_03_vol dd_21_vol, 2.13T
.
.
.
Cluster, Views, Exported capacity (GiB), Exported volumes
cluster-1, 2, 2209.47, 28
cluster-2, 1, 2178.46, 25
```

**See also**

- ◆ [report capacity-clusters on page 390](#)
- ◆ [report capacity-arrays on page 387](#)

## report create-monitors

Creates three performance monitors for each director in the VPLEX: storage-volume performance, port performance, and virtual volume performance. Each monitor has one file sink.

**Contexts** All contexts.

**Syntax**

```
report create-monitors
  [-d|--directory] directory
  [--force]
```

### Optional arguments

**[-d | -directory] *directory***- Directory in which to create the csv files.

Default directory path: /var/log/VPLEX/cli/reports/ on the management server.

**--force** - Forces the creation of the monitor, even if existing monitors are delayed in their polling.

**Description** Creates three monitors for each director in the VPLEX. Monitors are named:

- ◆ Cluster\_*n*\_Dir\_*nm*\_diskReportMonitor
- ◆ Cluster\_*n*\_Dir\_*nm*\_portReportMonitor
- ◆ Cluster\_*n*\_Dir\_*nm*\_volumeReportMonitor

The **period** attribute for the new monitors is set to 0 (automatic polling is disabled). Use the “[report poll-monitors](#)” command to force a poll.

Each monitor has one file sink. The file sinks are enabled.

By default, output files are located in /var/log/VPLEX/cli/reports/ on the management server. Output filenames are in the following format:

```
<Monitor name>_<Cluster_n_Dir_nm.csv
```

Disk report monitors collect:

- ◆ storage-volume.per-storage-volume-read-latency
- ◆ storage-volume.per-storage-volume-write-latency.

Port report monitors collect:

- ◆ be-prt.read
- ◆ be-prt.write
- ◆ fe-prt.ops
- ◆ fe-prt.read
- ◆ fe-prt.write

Volume report monitors collect:

- ◆ virtual-volume.ops
- ◆ virtual-volume.read
- ◆ virtual-volume.write

**Example** In the following example:

- ◆ The **report create-monitors** command creates a diskReportMonitor, portReportMonitor, and volumeReportMonitor for each director in the VPLEX,

- ◆ The `ll /monitoring/directors/*/monitors` command displays the new monitors:

```

VPlexcli:/> report create-monitors
Creating monitor diskReportMonitor on Director Cluster_1_Dir1A monitoring 981 targets, file
/var/log/VPlex/cli/reports/diskReportMonitor_Cluster_1_Dir1A.csv.
Successfully created 1 monitor(s) out of 1.

time: 11 sec
Creating monitor volumeReportMonitor on Director Cluster_1_Dir1A monitoring 30 targets, file
/var/log/VPlex/cli/reports/volumeReportMonitor_Cluster_1_Dir1A.csv.
Successfully created 1 monitor(s) out of 1.

time: 0 sec
Creating monitor portReportMonitor on Director Cluster_1_Dir1A monitoring 16 targets, file
/var/log/VPlex/cli/reports/portReportMonitor_Cluster_1_Dir1A.csv.
Successfully created 1 monitor(s) out of 1.
Creating monitor diskReportMonitor on Director
Cluster_1_Dir1A monitoring 981 targets, file
/var/log/VPlex/cli/reports/diskReportMonitor_Cluster_1_Dir1A.csv.
Successfully created 1 monitor(s) out of 1.
.
.
.
VPlexcli:/> ll /monitoring/directors/*/monitors

```

```

/monitoring/directors/Cluster_1_Dir1A/monitors:
Name Ownership Collecting Period Average Idle For Bucket Bucket Bucket Bucket
----- Ownership Data ----- Period ----- Min ----- Max ----- Width ----- Count
Cluster_1_Dir1A_diskReportMonitor true true 0s - 7.1min 100 1600100 25000 64
Cluster_1_Dir1A_portReportMonitor true true 0s - 6.88min - - - 64
Cluster_1_Dir1A_volumeReportMonitor true true 0s - 6.9min - - - 64

/monitoring/directors/Cluster_1_Dir1B/monitors:
Name Ownership Collecting Period Average Idle For Bucket Bucket Bucket Bucket
----- Ownership Data ----- Period ----- Min ----- Max ----- Width ----- Count
Cluster_1_Dir1B_diskReportMonitor true true 0s - 6.88min 100 1600100 25000 64
Cluster_1_Dir1B_portReportMonitor true true 0s - 6.68min - - - 64
Cluster_1_Dir1B_volumeReportMonitor true true 0s - 6.7min - - - 64
.
.
.

```

In the following example, the `--force` argument forces the creation of monitors, even though the creation results in missed polling periods:

```

VPlexcli:/> report create-monitors --force
Creating monitor diskReportMonitor on Director Cluster_1_Dir1A monitoring 981 targets, file
/var/log/VPlex/cli/reports/diskReportMonitor_Cluster_1_Dir1A.csv.
WARNING: One or more of your monitors is currently at least 25.0% behind its polling period.

Successfully created 1 monitor(s) out of 1.

WARNING: One or more of your monitors is currently at least 25.0% behind its polling period.

time: 1 sec
Creating monitor volumeReportMonitor on Director Cluster_1_Dir1A monitoring 30 targets, file
/var/log/VPlex/cli/reports/volumeReportMonitor_Cluster_1_Dir1A.csv.
WARNING: One or more of your monitors is currently at least 25.0% behind its polling period.

Successfully created 1 monitor(s) out of 1.

WARNING: One or more of your monitors is currently at least 25.0% behind its polling period.
.
.
.

```

- See also**
- ◆ [monitor add-file-sink on page 317](#)
  - ◆ [monitor create on page 320](#)
  - ◆ [monitor destroy on page 323](#)

- ◆ [monitor remove-sink on page 324](#)
- ◆ [report poll-monitors on page 395](#)

---

## report poll-monitors

Polls the report monitors created by the **report create-monitors** command.

**Contexts** All contexts.

**Syntax** `report poll-monitors`

**Description** The monitors created by the **report create-monitors** command have their **period** attribute set to 0 seconds (automatic polling is disabled) and one file sink.

Use this command to force an immediate poll and collection of performance data for monitors created by the **report create-monitors** command.

Output is written to files located in `/var/log/VPlex/cli/reports/` on the management server.

### Example

```
VPlexcli: /> report poll-monitors  
Collecting data for director Cluster_2_Dir_1B monitor Cluster_2_Dir_1B_diskReportMonitor.  
Collecting data for director Cluster_2_Dir_1B monitor Cluster_2_Dir_1B_portReportMonitor.  
Collecting data for director Cluster_2_Dir_1B monitor Cluster_2_Dir_1B_volumeReportMonitor.  
.  
.  
.
```

**See also**

- ◆ [monitor collect on page 319](#)
- ◆ [report create-monitors on page 392](#)

## rp import-certificate

Imports a RecoverPoint security certificate from the specified RPA cluster.

|                  |                       |
|------------------|-----------------------|
| <b>Contexts</b>  | All contexts.         |
| <b>Syntax</b>    | rp import-certificate |
| <b>Arguments</b> | None.                 |

**Description** This command runs an interview script to import the RecoverPoint security certificate.

In Metro systems, run this command on both management servers.



### **IMPORTANT**

**This command restarts the VPLEX CLI. All CLI and GUI sessions on the VPLEX cluster where the RPA cluster is attached.**

### **Before you begin**

You will need the IP address of the RecoverPoint cluster from which to import the security certificate.

**Example** Import the RecoverPoint security certificate from the RPA cluster at IP address 10.6.210.85:

```
VPlexcli:/> rp import-certificate  
This command will cause the VPLEX CLI process to restart if security settings are modified.  
This will require a new log in from all connected CLI and GUI clients.
```

```
To proceed type CONTINUE or hit enter to abort: CONTINUE
```

```
Please enter the IP v4 address of the RP cluster: 10.6.210.85
```

```
-----Certificate Details-----
```

```
Owner: CN=RecoverPoint, OU=Unified Storage Division, O=EMC Corporation, L=Ramat-Gan, ST=Israel, C=IL  
Issuer: CN=RecoverPoint, OU=Unified Storage Division, O=EMC Corporation, L=Ramat-Gan, ST=Israel, C=IL  
Serial number: 4d907d4c  
Valid from: Mon Mar 28 12:21:32 UTC 2011 until: Thu Mar 25 12:21:32 UTC 2021  
Certificate fingerprints:  
    MD5: CF:38:C3:55:A9:99:AC:A6:79:12:7C:83:C3:95:23:CB  
    SHA1: 4D:D6:29:30:ED:0A:77:6D:38:4E:10:D3:2E:37:29:CB:45:DC:9E:C0  
Signature algorithm name: SHA1withRSA  
Version: 3
```

```
Trust this certificate? (Y/N): Y
```

```
The management server console process will now restart, please press any key when you are ready.  
Please wait a minute before reconnecting.
```

```
Press '<Enter>' to continue ...<Enter>
```

```
Stopping EMC Vplex Management Console: Connection closed by foreign host.  
service@sms-advil-2:/opt/emc/VPlex/tools/utils>
```

- See also**
- ◆ [rp import-certificate on page 396](#)
  - ◆ [rp summary on page 403](#)
  - ◆ [rp validate-configuration on page 405](#)

## rp rpa-cluster add

Associates a cluster of RecoverPoint Appliances to single VPLEX cluster.

**Contexts** All contexts.

**Syntax**

```
rp rpa-cluster add
[-o|--host] rpa site management IP address
[-u|--admin-username] rpa admin-username
[-c|--cluster] cluster-id
```

**Arguments** **Required arguments**

**[-o|--host] rpa site management IP address** - \* The RPA cluster management IP address.

**[-u|--admin-username] admin-user-name** - \* The administrative username of the RPA

**Optional arguments**

**[-c|--cluster] cluster-id** - Context path of the VPLEX cluster associated with this cluster of RPAs. If no VPLEX cluster is specified, the ID of the local cluster is used. The local cluster is the cluster whose cluster-id matches the management server's IP seed. See [About cluster IP seed and cluster ID on page 434](#).

\* argument is positional.

**Description**

Adds information about a RecoverPoint Appliance cluster to VPLEX. Used by VPLEX to connect to RecoverPoint and retrieve replication information.

In Metro systems, run this command on both management servers.

---

**Note:** This command prompts for the RPA administrative password.

---

**Note:** Configuration of RPAs is not permitted during VPLEX NDU.

**Example** Add a RecoverPoint RPA cluster:

```
Vplexcli:/> rp rpa-cluster add -o 10.6.210.79 -u admin -c cluster-1
Enter rpa-cluster administrative password: Admin-password
Enter rpa-cluster administrative password again for verification:
Admin-password
```

**Example** Display RPA clusters.

- ◆ **ll /recoverpoint/rpa-clusters** displays summarized information about the RPA cluster.
- ◆ **ll /recoverpoint/rpa-cluster /ip-address** displays detailed information about the RPA cluster
- ◆ **ll /recoverpoint/rpa-cluster /ip-address/volumes/** displays the replication and journal volumes managed by the RPA
- ◆ **ls /recoverpoint/rpa-cluster /ip-address/volumes/volume-name** displays detailed information about the specified volume

```
Vplexcli:/> ll /recoverpoint/rpa-clusters
```

```
/recoverpoint/rpa-clusters:
RPA Host      VPLEX Cluster  RPA Site  RPA ID  RPA Version
-----
10.6.210.75  cluster-1      Site1     RPA 1   3.5(1.26)
```

```
VPLEXcli:/> 11 /recoverpoint/rpa-clusters/10.6.210.75
```

```
/recoverpoint/rpa-clusters/10.6.210.75:
```

Attributes:

| Name                   | Value       |
|------------------------|-------------|
| admin-username         | admin       |
| config-changes-allowed | true        |
| rp-health-indications  | []          |
| rp-health-status       | ok          |
| rp-software-serial-id  | -           |
| rpa-host               | 10.6.210.75 |
| rpa-id                 | RPA 1       |
| rpa-site               | Site1       |
| rpa-version            | 3.5(1.26)   |
| vplex-cluster          | cluster-1   |

Contexts:

| Name    | Description |
|---------|-------------|
| volumes | -           |

```
VPLEXcli:/> 11 /recoverpoint/rpa-clusters/10.6.210.75/volumes
```

```
/recoverpoint/rpa-clusters/10.6.210.75/volumes:
```

| Name                   | RPA Site | RP Type     | RP Role           | RP Group | VPLEX Group | Capacity |
|------------------------|----------|-------------|-------------------|----------|-------------|----------|
| RP_Repo_Vol_vol        | Site1    | Repository  | -                 | -        | RPTestCg    | 12G      |
| demo_DR1_vol           | Site1    | Replication | Production Source | cg1      | DemoSyncCG  | 5G       |
| demo_Replica_vol       | Site1    | Replication | Local Replica     | cg1      | Demo_Repli  | 5G       |
| demo_prodjournal_1_vol | Site1    | Journal     | -                 | cg1      | RPTestCg    | 5G       |
| demo_prodjournal_2_vol | Site1    | Journal     | -                 | cg1      | RPTestCg    | 5G       |
| demo_prodjournal_3_vol | Site1    | Journal     | -                 | cg1      | RPTestCg    | 5G       |
| .                      |          |             |                   |          |             |          |
| .                      |          |             |                   |          |             |          |
| .                      |          |             |                   |          |             |          |

```
VPLEXcli:/> 11 /recoverpoint/rpa-clusters/10.6.210.75/volumes/demo_DR1_vol
```

```
/recoverpoint/rpa-clusters/10.6.210.75/volumes/demo_DR1_vol:
```

| Name                      | Value                            |
|---------------------------|----------------------------------|
| rp-consistency-group      | cg1                              |
| rp-consistency-group-copy | prod                             |
| rp-replication-set        | RSet 1                           |
| rp-role                   | Production Source                |
| rp-type                   | Replication                      |
| rpa-site                  | Site1                            |
| size                      | 5G                               |
| uid                       | 600014400000010e03ec4a5dffffd8c3 |
| vplex-cluster             | cluster-1                        |
| vplex-consistency-group   | DemoSyncCG                       |

**Table 19 recoverpoint display fields (1 of 2)**

| Field                                                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>in /recoverpoint/rpa-clusters and rpa-clusters/ip-address contexts</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| admin-username                                                            | A pre-configured RecoverPoint user with an admin role granted all system permissions to manage RecoverPoint. Excluded privileges: downloading objects located on the RPAs, changing users and roles, security levels, and LDAP configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| config-changes-allowed                                                    | Whether or not RecoverPoint appliance configuration changes are currently allowed (maintenance mode). RecoverPoint appliance configuration changes are not allowed during VPLEX NDU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| rp-health-indications                                                     | If rp-health-status is anything other than OK, additional information about the problem and the component that is impacted.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| rp-health-status                                                          | Operational health of the RP cluster components including WAN, volumes, RPAs, and splitters.<br><b>OK</b> - All components of the RP configuration are operating as expected.<br><b>error</b> - One or more components of the RP configuration is not operating as expected. The <b>rp-health-indications</b> field displays additional information.<br><b>warning</b> - One or more components of the RP configuration is not operating as expected. The <b>rp-health-indications</b> field displays additional information.<br><b>unknown</b> - VPLEX cannot connect to the RPA cluster.                                                                                                                                                                                                                                             |
| RPA Host                                                                  | Management IP address of the RPA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| VPLEX Cluster                                                             | Cluster ID or name of the VPLEX cluster associated with the RPA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RPA Site                                                                  | Name of the RPA site. There can be up to two sites in a RecoverPoint installation; a local site and a remote site. In one-site configurations (CDP), both the production and local copy reside at the local site. In two-site configurations (stretch CDP, CRR, and CLR), the production copy is at the local site and the remote copy is at the remote site.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| RPA ID                                                                    | ID of the primary/preferred RPA.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RPA Version                                                               | Version of RecoverPoint on the RPA, in the following format:<br>MajorVersion.MinorVersion.ServicePack.Patch (branch.build)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>in /recoverpoint/rpa-clusters/ip-address/volumes context</b>           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Name                                                                      | Name of the volume as it appears in VPLEX.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| RPA Site                                                                  | See <a href="#">"RPA Site"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| RP Type                                                                   | Role the volume is serving in RecoverPoint replication.<br><b>Journal</b> - One or more volumes dedicated on the storage at each copy in a RecoverPoint configuration. Journals are defined per copy, and can consist of multiple journal volumes.<br><b>Replication</b> - All volumes of a replication set. One production volume and one or two replica volumes.<br><b>Repository</b> - A special volume that must be dedicated on the SAN-attached storage at each site, for each RPA cluster. The repository volume serves all RPAs of the particular cluster and splitters associated with that cluster. It stores configuration information about the RPAs and RecoverPoint consistency groups, which enables a properly functioning RPA to seamlessly assume the replication activities of a failing RPA from the same cluster. |

**Table 19 recoverpoint display fields (2 of 2)**

| Field                                                                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RP Role                                                                | <p>The function of a copy, in regards to RecoverPoint failover. Initially, RecoverPoint configurations can consist of a Production Source and a local replica and/or a remote replica.</p> <ul style="list-style-type: none"> <li>• After a temporary failover, the Production Source becomes the Replica at Production and the Remote Replica (or Local Replica) becomes the Remote (or Local) Source.</li> <li>• After a permanent failover: the Replica at Production becomes the Production Source and the Remote (or Local) Source becomes the Remote (or Local) Replica.</li> </ul> <p><b>Local Replica</b> - The replica at the local site. Also, the role of the local copy, before failover.<br/> <b>Production Source</b> - The role of a consistency group production copy, before failover.<br/> <b>Remote Replica</b> - The replica at the remote site that is being replicated to in CRR or CLR configurations. Also denotes the role of the replica at the remote site, before failover.</p> |
| RP Group                                                               | If the volume is a member of a RecoverPoint consistency group, the name of the group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| VPLEX Group                                                            | <p>The VPLEX consistency group to which this volume belongs. Production and replica volumes associated with RecoverPoint must be in VPLEX consistency groups that have the following attributes:</p> <ul style="list-style-type: none"> <li>• Cache-mode property is synchronous</li> <li>• Consistency groups with the “visibility” property set to both clusters must also have their “storage-at-clusters” set to both clusters.</li> <li>• Recoverpoint-enabled property set to true.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| Capacity                                                               | Capacity of the volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>In /recoverpoint/rpa-clusters/ip-address/volumes/volume context</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| rp-consistency-group                                                   | See “RP Group”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| rp-consistency-group-copy                                              | <p>The RecoverPoint copy to which the volume belongs. In CDP and CRR configurations, there is one production copy and one replica copy. In CLR configurations, there is one production copy and two replica copies (one local copy at the production site and one remote copy at the disaster recovery site). The production copy consists of production volumes and the production journal, which may consist of one or more journal volumes. The non-production copies (i.e. replica copies) each consist of replica volumes and a replica journal, which may consist of one or more journal volumes.</p>                                                                                                                                                                                                                                                                                                                                                                                                 |
| rp-replication-set                                                     | <p>The RecoverPoint replication set to which the volume belongs. Replication sets consist of the production source volume and the replica volume(s) to which it replicates. Every storage volume in the production storage must have a corresponding volume at each copy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| rp-role                                                                | See “RP Role”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| rp-type                                                                | See “RP Type”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| rpa-site                                                               | See “RPA Site”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| rpa-software-serial-d                                                  | Software serial ID for each rpa-cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| size                                                                   | See “Capacity”.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| uid                                                                    | Unique Identifier for the volume. A 64-bit number used to uniquely identify each VPLEX volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| vplex-cluster                                                          | The VPLEX cluster with which this RPA cluster is associated.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| vplex-consistency-group                                                | The name of the VPLEX consistency group of which this volume is a member.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

- See also**
- ◆ [rp rpa-cluster remove on page 402](#)
  - ◆ [rp summary on page 403](#)

- ◆ [rp validate-configuration](#) on page 405

## rp rpa-cluster remove

Removes information about a RecoverPoint Appliance from VPLEX.

**Contexts** All contexts.

**Syntax** `rp rpa-cluster remove  
[-r|--rpa] IP-address`

**Arguments** **Required arguments**  
`[-r|--rpa-cluster] IP-address` - The site management IP address of the RPA cluster to remove.

**Description** Removes information about a RPA cluster from VPLEX.  
Use the `ll` command in `/recoverpoint/rpa-clusters` context to display the site management IP address.

**Example** Remove an RPA:

```
Vplexcli:/> ll /recoverpoint/rpa-clusters

/recoverpoint/rpa-clusters:
RPA Host      VPLEX Cluster  RPA Site  RPA ID  RPA Version
-----
10.6.210.75   cluster-1      Site1     RPA 1   3.5(1.26)

Vplexcli:/> rp rpa-cluster remove --rpa-cluster 10.6.210.75
```

**See also**

- ◆ [rp rpa-cluster add on page 397](#)
- ◆ [rp summary on page 403](#)

## rp summary

Displays a summary of replication for the entire VPLEX cluster, across all connected RPA sites/clusters.

**Contexts** All contexts.

In /recoverpoint/ context, command is **summary**.

**Syntax** `rp summary`

**Arguments** **Required arguments**  
None.

**Description** This command calculates the total number of volumes and the total capacity for each RP type and RP role it finds in the /recoverpoint/rpa-clusters context.  
Also prints cumulative information.

**Example** Display a VPLEX Metro with RecoverPoint RPAs deployed at both VPLEX clusters:

```
VPlexcli:/> rp summary
```

```
VPlexcli:/> rp summary
```

```
RecoverPoint Replication Totals:
```

| VPLEX Cluster | RP Type            | RP Role           | Total Volumes | Total Capacity |
|---------------|--------------------|-------------------|---------------|----------------|
| cluster-1     | Replication        | Production Source | 2             | 10G            |
|               |                    | Local Replica     | 2             | 10G            |
|               |                    | Remote Replica    | 0             | 0G             |
|               | Journal Repository |                   | 4             | 20G            |
|               |                    |                   | 1             | 10G            |
|               | Totals:            |                   |               | 9              |

| VPLEX Cluster | RP Type            | RP Role           | Total Volumes | Total Capacity |
|---------------|--------------------|-------------------|---------------|----------------|
| cluster-2     | Replication        | Production Source | 1             | 20G            |
|               |                    | Local Replica     | 1             | 20G            |
|               |                    | Remote Replica    | 0             | 0G             |
|               | Journal Repository |                   | 1             | 20G            |
|               |                    |                   | 1             | 10G            |
|               | Totals:            |                   |               | 4              |

**Table 20** **rp summary display fields (1 of 2)**

| Field         | Description                                                   |
|---------------|---------------------------------------------------------------|
| VPLEX Cluster | Name of the VPLEX cluster associated with a RecoverPoint RPA. |

**Table 20** **rp summary display fields (2 of 2)**

| Field                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RP Type              | <p>Type of volume.</p> <p><b>Journal</b> - One or more volumes dedicated on the storage at each copy in a RecoverPoint configuration. Journals are defined per copy, and can consist of multiple journal volumes.</p> <p><b>Replication</b> - All volumes of a replication set. One production volume and one or two replica volumes.</p> <p><b>Repository</b> - A special volume that must be dedicated on the SAN-attached storage at each site, for each RPA cluster. The repository volume serves all RPAs of the particular cluster and splitters associated with that cluster. It stores configuration information about the RPAs and RecoverPoint consistency groups, which enables a properly functioning RPA to seamlessly assume the replication activities of a failing RPA from the same cluster.</p>                                                                                                                                                         |
| RP Role              | <p>The function of a copy, in regards to failover.</p> <p>Initially, RecoverPoint configurations can consist of a Production Source and a Local Replica and/or a Remote Replica.</p> <ul style="list-style-type: none"> <li>• After a temporary failover, the Production Source becomes the Replica at Production and the Remote Replica (or Local Replica) becomes the Remote (or Local) Source.</li> <li>• After a permanent failover: the Replica at Production becomes the Production Source and the Remote (or Local) Source becomes the Remote (or Local) Replica.</li> </ul> <p><b>Local Replica</b> - The replica at the local site. Also, the role of the local copy, before failover.</p> <p><b>Production Source</b> - Production copy, before failover.</p> <p><b>Remote Replica</b> - The replica at the remote site that is being replicated to in CRR or CLR configurations. Also denotes the role of the replica at the remote site, before failover.</p> |
| Total Number Volumes | Number of volumes protected by RecoverPoint at the specified VPLEX cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Total Capacity       | Total capacity of the volumes protected by RecoverPoint at the specified VPLEX cluster.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

**See also**

- ◆ [rp rpa-cluster add on page 397](#)
- ◆ [rp rpa-cluster remove on page 402](#)
- ◆ [rp validate-configuration on page 405](#)

## rp validate-configuration

Validates the RecoverPoint splitter configuration.

**Contexts** All contexts.

In /recoverpoint/ context, command is **validate-configuration**.

**Syntax** `rp validate-configuration`

**Arguments** **Required arguments**  
None.

**Description** This command checks the system configuration with respect to RecoverPoint and displays errors or warnings if errors are detected.  
For VPLEX Metro configurations, run this command on both management servers.



### CAUTION

When RPAs are zoned to VPLEX using single-channel mode (2 RPA ports are zoned to VPLEX front end ports, and 2 RPA ports are zoned to the VPLEX back end ports) this command reports the ports as "WARNING". This is because the command checks that all 4 ports on the RPA are zoned to both VPLEX front end *and* backend ports (dual-channel mode). See [Example on page 406](#).

**Best practice is to zone every RPA port to both VPLEX front end and backend ports. For configurations where this is not possible or desirable, this command detects the port configuration, displays a warning. Administrators who have purposely configured single-channel mode can safely ignore the warning.**

This command performs the following checks:

Splitter checks:

- ◆ VPLEX splitters are installed.
- ◆ All splitter versions agree.
- ◆ The VPLEX splitter status is OK.

RecoverPoint cluster checks:

- ◆ VPLEX management server can reach all the attached RecoverPoint clusters

Storage view checks:

- ◆ Storage views do not have mixed RecoverPoint and non-RecoverPoint initiator ports.
- ◆ RecoverPoint storage views have access to multiple ports.
- ◆ No volume is exposed to more than one RecoverPoint cluster.
- ◆ No RecoverPoint journal or repository volumes are exposed to hosts.

Initiator port checks:

- ◆ All the RecoverPoint initiator ports are registered.
- ◆ All the registered RecoverPoint initiator ports are used.

RP Cluster:

- ◆ VPLEX management server can reach all the attached RecoverPoint Clusters.

Consistency group checks:

- ◆ VPLEX consistency groups are aligned with RecoverPoint consistency groups.

Volumes checks:

- ◆ All production volumes are not remote volumes.
- ◆ All distributed production volumes have detach rule set correctly.
- ◆ All distributed production volumes have cache mode set correctly.
- ◆ All production and replica volumes are in RecoverPoint-enabled VPLEX consistency groups.
- ◆ No replica volumes is a remote volume.
- ◆ All distributed replica volumes have detach rule set correctly.
- ◆ All distributed replica volumes have cache mode set correctly.
- ◆ All journal and repository volumes are local volumes.
- ◆ All repository volumes are not in any VPLEX consistency group.

**Example** Check a healthy RecoverPoint configuration:

```
VPLexcli:/> rp validate-configuration
This command may take several minutes to complete. Please be patient.
=====
Validate the VPLEX Splitters
=====
Validating that VPLEX Splitters are installed
OK
Validating that all splitter versions agree
OK
Validating the VPLEX Slitter status
OK
=====
Validate Storage Views
=====
Validating that storage views do not have mixed non-recoverpoint and recoverpoint initiator
ports                                     OK
Validating that recoverpoint storageviews have access to multiple recoverpoint ports
OK
=====
Validate Initiator Ports
=====
Validating that all the recoverpoint initiator ports are registered
OK
Validating that all the registered recoverpoint initiator ports are used
OK
=====
Validation Summary
=====
Your system appears healthy. There were no problems found in the system.
```

**Example** Detect RPA ports that are zoned as single-channel mode:

```
VPLexcli:/> rp validate-configuration
This command may take several minutes to complete. Please be patient.
.
.
.
=====
Validate the VPLEX to RPA zoning
=====
```

Validating that VPLEX sees all expected initiator ports from RPAs  
WARNING

VPLEX does not see RPA initiator port: 0x500124804dc50283

VPLEX does not see RPA initiator port: 0x500124824dc50283

VPLEX does not see RPA initiator port: 0x500124804a00021b

VPLEX does not see RPA initiator port: 0x500124824a00021b

Validating that VPLEX sees all expected backend RPA ports  
WARNING

VPLEX does not see RPA back end port: 0x500124814dc50283

VPLEX does not see RPA back end port: 0x500124834dc50283

VPLEX does not see RPA back end port: 0x500124814a00021b

VPLEX does not see RPA back end port: 0x500124834a00021b

=====  
Validation Summary

=====  
The following potential problems were found in the system:

1 problem(s) were found with RecoverPoint Clusters.

8 potential problem(s) were found with the zoning between VPLEX and RecoverPoint.

### **Example** Detect a mis-configured storage-view:

```
VPlexcli:/> rp validate-configuration
```

```
.  
.  
.
```

Validating that storage views do not have mixed non-recoverpoint and recoverpoint initiator ports

ERROR

Storage view rp-view-demo has mixed types of initiator ports.

```
.  
.  
.
```

=====  
Validation Summary

=====  
The following potential problems were found in the system:

1 problem(s) were found with storage views.

- See also**
- ◆ [rp rpa-cluster add on page 397](#)
  - ◆ [rp rpa-cluster remove on page 402](#)
  - ◆ [rp summary on page 403](#)



---

## schedule add

Schedules a job to run at the specified time(s).

**Contexts** All contexts.

**Syntax**  
`schedule add`  
`[-t|--time] time`  
`[-c|--command] command`

**Arguments** **Required arguments**

`[-t|--time] time` - \* Date and time the job executes in crontab-style format enclosed in quote marks. Values are specified in the crontab-style format:

“<Minute> <Hour> <Day of the Month> <Month> <Day of the week>”

**Minute** - 0-59.

**Hour** - 0-23.

**Day of the Month** - 1-31.

**Month** - 1-12, January = 1...December = 12

**Day of the week** - 0-6, Sunday = 0...Saturday = 6

`[-c|--command] command` - \* The CLI command to be executed at the specified time.

\* - argument is positional.

**Example** To run the **tree** command every night at 1:00 a.m.:

```
Vplexcli:/> schedule add " * 1 * * *" --command tree
```

**See also**

- ◆ [schedule list on page 410](#)
- ◆ [schedule modify on page 411](#)
- ◆ [schedule remove on page 412](#)

---

## schedule list

Lists all scheduled jobs.

**Contexts** All contexts.

**Syntax** `schedule list`

**Example**

```
Vplexcli:/> schedule list
[0] 30 13 * * 3 syrcollect
[1] * 1 * * * tree
[2] * 2 * * * tree
```

**See also**

- ◆ [schedule add on page 409](#)
- ◆ [schedule modify on page 411](#)
- ◆ [schedule remove on page 412](#)

## schedule modify

Modifies an existing scheduled job.

**Contexts** All contexts.

**Syntax**

```
schedule modify
[-j|--job] job ID
[-t|--time] time
[-c|--command] command
```

**Arguments** **Required arguments**

**[-j|--job] job ID** - \* ID of the scheduled job as displayed by the “[schedule list](#)” command.

**[-t|--time] time** - \* Date and time the job executes in crontab-style format enclosed in quote marks. Values are specified in the following format:

“<Minute> <Hour> <Day of the Month> <Month> <Day of the week>”

**Minute** - 0-59.

**Hour** - 0-23.

**Day of the Month** - 1-31.

**Month** - 1-12, January = 1...December = 12

**Day of the week** - 0-6, Sunday = 0...Saturday = 6

**[-c|--command] command** - \* The CLI command to be executed at the specified time.

\* - argument is positional.

**Example** To modify a job with the ID of 3 so that it runs every day at 11:00 a.m. type:

```
Vplexcli:/> schedule list
[0] 30 13 * * 3 syrcollect
[1] * 1 * * * tree
[2] * 2 * * * tree
[3] * 3 * * * tree
```

```
Vplexcli:/> schedule modify 3 -t "*" 11 * * *" -c tree
```

**See also**

- ◆ [schedule list on page 410](#)
- ◆ [schedule remove on page 412](#)

## schedule remove

Removes a scheduled job.

**Contexts** All contexts.

**Syntax** `schedule remove`  
`[-j|--job] job ID`

**Arguments** **Required arguments**  
`[-j|--job] job ID` - \* ID of the scheduled job as displayed by the “[schedule list](#)” command.

\* - argument is positional.

**Example** Remove job with the ID of 3:

```
Vplexcli:/> schedule list  
[0] 30 13 * * 3 syrcollect  
[1] * 1 * * * tree  
[2] * 2 * * * tree  
[3] * 3 * * * tree
```

```
Vplexcli:/> schedule remove 3  
Removed scheduled job 3.
```

```
Vplexcli:/> schedule list  
[0] 30 13 * * 3 syrcollect  
[1] * 1 * * * tree  
[2] * 2 * * * tree
```

**See also**

- ◆ [schedule list on page 410](#)
- ◆ [schedule modify on page 411](#)

## scheduleSYR add

Schedules a weekly SYR data collection.

**Contexts** All contexts.

**Syntax**

```
scheduleSYR add
  [-d|--dayOfWeek] 0-6
  [-t|--hours] 0-23
  [-m|--minutes] 0-59
```

**Arguments**

**Required arguments**

**[-d|--dayOfWeek] 0-6** - Day of the week run the collection.  
Valid values are 0-6, where Sunday = 0...Saturday = 6.

**[-t|--hours] 0-23** - Hour at which to run the collection.

**[-m|--minutes] 0-59** - Minute at which to run the collection.

**Description** Typically, SYR collection and reporting are configured at initial system setup. Use this command to add a scheduled SYR collection time if none was configured.

SYR data collection can be scheduled to occur at most once a week. Attempts to add another weekly schedule results in an error.

SYR reporting gathers VPLEX configuration files and forward them to EMC. SYR reports provide:

- ◆ Faster problem resolution and RCA
- ◆ Proactive maintenance
- ◆ Data for performance analysis

To modify the existing SYR collection time, use the “[scheduleSYR remove](#)” command to delete the current time, and the “[scheduleSYR add](#)” command to specify a new collection time.

**Example** Schedule a SYR collection for every Monday at 11:30 p.m.:

```
Vplexcli:/> scheduleSYR add -d 3 -t 12 -m 30
SYR data collection job scheduled
```

```
Vplexcli:/> scheduleSYR list
SYR data collection job is currently scheduled at:
Day of Week: 3 (Sunday=0, Monday=1,...Saturday=6)
Hours: 12
Minutes: 30
```

**See also**

- ◆ [configuration event-notices-reports config](#) on page 121
- ◆ [configuration event-notices-reports reset](#) on page 124
- ◆ [schedule list](#) on page 410
- ◆ [scheduleSYR list](#) on page 414
- ◆ [scheduleSYR remove](#) on page 415
- ◆ [syrcollect](#) on page 482

## scheduleSYR list

Lists the scheduled SYR data collection job.

**Contexts** All contexts.

**Syntax** `scheduleSYR list`

**Example** List the SYC collection schedule:

```
VPlxcli:/> scheduleSYR list  
SYR data collection job is currently scheduled at:  
Day of Week: 1 (Sunday=0, Monday=1,...Saturday=6)  
Hours: 23  
Minutes: 30
```

**See also**

- ◆ [configuration event-notices-reports config on page 121](#)
- ◆ [configuration event-notices-reports reset on page 124](#)
- ◆ [scheduleSYR add on page 413](#)
- ◆ [scheduleSYR remove on page 415](#)

---

## scheduleSYR remove

Removes the currently scheduled SYR data collection job.

**Contexts** All contexts.

**Syntax** `scheduleSYR remove`

**Description** Only one SYR data collection can be scheduled. The current SYR collection cannot be modified. To modify the SYR data collection job:

- ◆ Use the **scheduleSYR remove** command to remove the existing collection job.
- ◆ Use the “[scheduleSYR add](#)” command to create a new collection job.

**Example** Remove a scheduled collection:

```
Vplexcli:/> scheduleSYR remove  
Removing SYR data collection job scheduled at:  
Day of Week: 3 (Sunday=0, Monday=1,...Saturday=6)  
Hours: 13  
Minutes: 30  
SYR data collection job removed successfully
```

- See also**
- ◆ [configuration event-notices-reports config on page 121](#)
  - ◆ [configuration event-notices-reports reset on page 124](#)
  - ◆ [scheduleSYR add on page 413](#)
  - ◆ [scheduleSYR list on page 414](#)

## script

Changes to interactive Jython scripting mode.

**Contexts** All contexts.

**Syntax**

```
script
  [-i|--import] module
  [-u|--unimport] module
```

### Optional arguments

**[-i|--import] *module*** - Import the specified Jython module without changing to interactive mode. After importation, commands registered by the module are available in the CLI. If the module is already imported, it is explicitly reloaded.

**[-u|--unimport] *module*** - Unimport the specified Jython module without changing to interactive mode. All the commands that were registered by that module are unregistered.

**Description** Changes the command mode from VPLEX CLI to Jython interactive mode.

To return to the normal CLI shell, type a period '.' and press **ENTER**.

Use the **--import** and **--export** arguments to import/export the specified Jython module without changing to interactive mode.

**Example** Enter Jython interactive mode:

```
Vplexcli:/> script
Jython 2.2 on java1.6.0_03

>>>
```

Exit Jython interactive mode:

```
>>> .
```

```
Vplexcli:/>
```

Import/unimport the specified Jython module without changing to interactive mode:

```
Vplexcli:/> script --import ndu
```

```
Vplexcli:/> script --unimport ndu
```

**See also** ♦ [source on page 455](#)

## security create-ca-certificate

Creates a new Certification Authority (CA) certificate.

**Contexts** All contexts.

**Syntax**

```
security create-ca-certificate
[-l|--keylength] 384-2048
[-d|--days] 365-1825
[-o|--ca-cert-outfilename] filename
[-f|--ca-key-outfilename] filename
[-s|--ca-subject-filename] filename
```

### Optional arguments

**[-l|--keylength] length** - The length (number of bits) for the CA key.

Default: 2048. Range: 384 - 2048

**[-d|--days] days** - Number of days that the certificate is valid.

Default: 1825 (5 years). Range: 365 - 1825.

**[-o|--ca-cert-outfilename] filename** - CA Certificate output filename.

Default: **strongswanCert.pem**.

**[-f|--ca-key-outfilename] filename** - CA Key output filename.

Default: **strongswanKey.pem**.

**[-s|--ca-subject-filename] filename** - Name of the CA subject information file that contains the subject information to create the CA certificate.

**Description** A management server authenticates users against account information kept on its local file system. An authenticated user can manage resources in the cluster(s).

The Certification Authority (CA) is used to sign management server certificates.

The **security create-ca-certificate** and **security create-host-certificate** commands create the CA and host certificates using a pre-configured Distinguished Name where the Common Name is the VPLEX cluster Top Level Administrator (TLA). If the TLA is not already set, it must be set manually to prevent certificate creation failure.

Alternatively, use the **--ca-subject-filename** argument to create a custom Distinguished Name. Specify the full path of the subject file unless the subject file is in the local CLI directory.

This command creates two objects on the management server:

- ◆ A CA certificate file valid for 1825 days (5 years). The file is located at:  
/etc/ipsec.d/cacerts/strongswanCert.pem
- ◆ A private key protected by a passphrase. The CA key is located at:  
/etc/ipsec.d/private/strongswanKey.pem

### Using the security commands

**Note:** Take note of the passphrases you use to create these certificates and save them in a secure location. They will be required at other times when maintaining the VPLEX clusters.

The steps to configure security on a VPLEX differ depending on whether the configuration is a VPLEX Local or VPLEX Metro/Geo.

For VPLEX Local configurations:

1. Use the **security create-ca-certificate** command to create the CA certificate.
2. Write down the passphrase entered during the creation of CA certificate.
3. Use the **security create-host-certificate** command to create the host certificate.  
A prompt is displayed for the passphrase used to create the CA certificate.
4. Write down the passphrase entered during the creation of host certificate.
5. Use the **security ipsec-configure** command to start the VPN process.  
A prompt for the passphrase used to create the host certificate is displayed.

For VPLEX Metro and Geo configurations:

**Note:** Take note of the passphrases you use to create these certificates and save them in a secure location. They will be required at other times when maintaining the VPLEX clusters.

On the first cluster:

1. **scp** the CA certificate file (default name: **strongswanCert.pem**)  
from: **/etc/ipsec.d/cacerts** on the first cluster  
to: **/etc/ipsec.d/cacerts** directory on the second cluster
2. **scp** the CA passphrase (default name: **strongswanKey.pem**)  
from: **/etc/ipsec.d/private** directory on the first cluster  
to: **/etc/ipsec.d/private** directory on the second cluster

For example:

```
scp /etc/ipsec.d/cacerts/strongswanCert.pem
  service@10.12.177.74:/etc/ipsec.d/cacerts

scp /etc/ipsec.d/private/strongswanKey.pem
  service@10.12.177.74:/etc/ipsec.d/private
```

On the second cluster:

3. Use the **security create-host-certificate** command to create the host certificate.  
A prompt is displayed for the passphrase used to create the CA certificate.
4. Write down the passphrase entered during the creation of host certificate.

Set up the VPN between clusters:

On the first cluster:

5. Use the **security ipsec-configure** command to start the VPN process.  
A prompt for the passphrase used to create the host certificate is displayed.

```
security ipsec-configure -l 1 -r 2 -i 10.12.177.74
```

On the second cluster:

6. Use the **security ipsec-configure** command to start the VPN process.  
A prompt for the passphrase used to create the host certificate on the second cluster is displayed.

```
security ipsec-configure -l 2 -r 1 -i 10.12.177.71
```



## IMPORTANT

**In VPLEX Metro and Geo configurations: Use the same CA certificate passphrase for host certificates on management servers 1 and 2.**

**Example** Create a default CA certificate with the default CA certificate subject information.

```
Vplexcli:/> security create-ca-certificate
```

```
Please enter the passphrase for the CA Certificate Key:
```

```
Re-enter the passphrase for the Certificate Key:  
New CA certificate strongswanCert.pem created  
New CA key strongswanKey.pem created
```

**Example** Create a default CA certificate (**strongswanCert.pem**) with custom CA certificate subject information. In the following example,

- The “`security create-certificate-subject`” command creates a custom subject file named `TestSubject.txt`.
- The `security create-ca-certificate` command creates a default CA certificate with the specified custom subject file

```
Vplexcli:/> security create-certificate-subject -c US -s NewYork -m EMC -u EMC -l NewYork -n  
CommonTestName -e test@emc.com -o TestSubject.txt
```

```
Vplexcli:/> security create-ca-certificate --ca-subject-filename TestSubject.txt
```

The following files are created in the specified directories:

- `/etc/ipsec.d/cacerts/strongswanCert.pem`
- `/etc/ipsec.d/private/strongswanKey.pem`

**Example** Create a custom CA certificate with a custom CA certificate subject information. In the following example:

- The `security create-certificate-subject` command creates a custom subject file named `TestSubject.txt`.
- The `security create-ca-certificate` command creates a custom CA certificate with the specified custom subject file.

```
Vplexcli:/> security create-certificate-subject -c US -s NewYork -m EMC -u EMC -l NewYork -n  
CommonTestName -e test@emc.com -o TestSubject.txt
```

```
Vplexcli:/> security create-ca-certificate --ca-cert-outfilename TestCACert.pem  
--ca-key-outfilename TestCAKey.pem --ca-subject-filename TestSubject.txt
```

The following files are created in the specified directories:

- `/etc/ipsec.d/cacerts/TestCACert.pem`
- `/etc/ipsec.d/private/TestCAKey.pem`

- See also**
- ◆ [security create-certificate-subject on page 421](#)
  - ◆ [security create-host-certificate on page 423](#)
  - ◆ [security delete-ca-certificate on page 425](#)
  - ◆ [security delete-host-certificate on page 426](#)
  - ◆ [security export-ca-certificate on page 427](#)
  - ◆ [security export-host-certificate on page 428](#)

- ◆ [security import-ca-certificate](#) on page 429
- ◆ [security import-host-certificate](#) on page 431
- ◆ [security ipsec-configure](#) on page 433
- ◆ [security show-cert-subj](#) on page 442
- ◆ *EMC VPLEX Security Configuration Guide*

## security create-certificate-subject

Creates a subject file used in creating security certificates.

**Contexts** All contexts.

**Syntax**

```
security create-certificate-subject
[-c|--country] country
[-s|--state] state
[-m|--org-name] organizational name
[-u|--org-unit] organizational unit
[-l|--locality] locality
[-n|--common-name] name
[-e|--email] e-mail
[-o|--subject-out-filename] filename
[--force]
```

**Arguments** **Required arguments**

**[-o|--subject-out-filename] *filename*** - The filename of the subject file.

**Optional arguments**

**[-c|--country] *country*** - The Country value for the country key in the subject file.

**[-s|--state] *state*** - The State value for the state key in the subject file.

**[-m|--org-name] *organizational name*** - Organizational Name value for the organizational name key in the subject file.

**[-u|--org-unit] *organizational unit*** - Organizational Unit value for the organizational unit key in the subject file.

**[-l|--locality] *locality*** - Locality value for the locality key in the subject file.

**[-n|--common-name] *name*** - Name value for the name key in the subject file.

**[-e|--email] *e-mail*** - E-mail value for the e-mail key in the subject file.

**[-o|--subject-out-filename] *Output Subject Filename*** - The filename of the subject file to be created.

**--force** - Overwrites the specified subject-out-filename if a file of that name already exists. If a file with the subject-out-filename already exists and the **--force** argument is not specified, the command fails.

**Description** Creates a subject file used in certificate creation.

**Example** Create a default certificate subject file:

```
VPlexcli: /> security create-certificate-subject --subject-out-filename TestSubject.txt
The certificate subject file TestSubject.txt file is created at /var/log/VPlex/cli
directory with the following information:
```

```
SUBJECT_COUNTRY=US
SUBJECT_STATE=Massachusetts
SUBJECT_LOCALITY=Hopkinton
SUBJECT_ORG=EMC
SUBJECT_ORG_UNIT=EMC
SUBJECT_COMMON_NAME=FNMM00102200421
SUBJECT_EMAIL=support@emc.com
```

**Example** Create a custom certificate subject file:

```
Vplexcli:/> security create-certificate-subject -c US -s NewYork -m EMC -u EMC -l NewYork -n  
CommonTestName -e test@emc.com --subject-out-filename TestSubject.txt
```

The certificate subject file **TestSubject.txt** file is created at `/var/log/Vplex/cli` directory with the following information:

```
SUBJECT_COUNTRY=US  
SUBJECT_STATE=NewYork  
SUBJECT_LOCALITY=NewYork  
SUBJECT_ORG=EMC  
SUBJECT_ORG_UNIT=EMC  
SUBJECT_COMMON_NAME=CommonTestName  
SUBJECT_EMAIL=test@emc.com
```

- See also**
- ◆ [security create-ca-certificate on page 417](#)
  - ◆ [security create-host-certificate on page 423](#)
  - ◆ [security export-ca-certificate on page 427](#)
  - ◆ [security export-host-certificate on page 428](#)
  - ◆ [security import-ca-certificate on page 429](#)
  - ◆ [security import-host-certificate on page 431](#)
  - ◆ [security ipsec-configure on page 433](#)
  - ◆ [security show-cert-subj on page 442](#)
  - ◆ *EMC VPLEX Security Configuration Guide*

## security create-host-certificate

Creates a new host certificate and signs it with an existing CA certificate.

**Contexts** All contexts.

**Syntax**

```
security create-host-certificate
[-l|--keylength] 384-2048
[-d|--days] 365-730
[-o|--host-cert-outfilename] filename
[-f|--host-key-outfilename] filename
[-r|--ca-subject-filename] filename
[-c|--ca-cert-filename] ca certificate
[-k|--ca-key-filename] ca key
[-s|--host-subject-filename] filename
[-g|--get-master-ca]
```

### Optional arguments

**[-l|--keylength] length** - The length (number of bits) for the CA key.

Default: 2048. Range: 384 - 2048.

**[-d|--days] days** - Number of days that the certificate is valid.

Default: 730 (2 years). Range: 365 - 730.

**[-o|--host-cert-outfilename] filename** - Host certificate output filename.

Default: **hostCert.pem**.

**[-f|--host-key-outfilename] filename** - Host key output filename.

Default: **hostKey.pem**.

**[-r|--host-cert-req-outfilename] filename** - host certificate Request output filename.

Default: **hostCertReq.pem**.

**[-c|--ca-cert-filename] ca certificate** - CA certificate used to sign the host certificate.

Default: **strongswanCert.pem**.

**[-k|--ca-key-filename] ca key** - CA Key used to sign the host certificate.

Default: **strongswanKey.pem**

**[-s|--host-subject-filename] filename** - File that contains the subject information to create the host certificate.

**[-g|--get-master-ca]** - Pulls the master CA to the requesting cluster and creates the digital certificate. Copies the updated serial file back to the master server so that the master server has a serial number that is up to date with the number of digital certificates that the CA created. Enables strict browsers (FireFox) to connect to different clusters from the same browser.

**Description** Generates a host certificate request and signs it with the Certification Authority certificate created by the **security create-ca-certificate** command.

The CA Certificate and CA Key must be created prior to running this command.

The host certificate is stored at `/etc/ipsec.d/certs`.

The host key is stored at `/etc/ipsec.d/private`.

The host certificate request is stored at `/etc/ipsec.d/reqs`.

The CA certificate file is read from /etc/ipsec.d/cacerts.

The CA Key is read from /etc/ipsec.d/private.

**Example** Create a default host certificate with the default host certificate subject information:

```
VPllexcli:/> security create-host-certificate
```

The following files are created in the specified directories:

- /etc/ipsec.d/certs/hostCert.pem
- /etc/ipsec.d/private/hostKey.pem

**Example** Create a custom host certificate with the default host certificate subject information:

```
VPllexcli:/> security create-host-certificate --host-cert-outfilename  
TestHostCert.pem --host-key-outfilename TestHostKey.pem
```

The following files are created in the specified directories:

- /etc/ipsec.d/certs/TesthostCert.pem
- /etc/ipsec.d/private/TesthostKey.pem

**Example** Create a custom host certificate with custom host certificate subject information. In the following example:

- The **security create-certificate-subject** command creates a custom subject file named TestSubject.txt.
- The **security create-host-certificate** command creates a custom host certificate with the specified custom subject file.

```
VPllexcli:/> security create-certificate-subject -c US -s NewYork -m EMC -u EMC -l NewYork -n  
CommonTestName -e test@emc.com -o TestSubject.txt
```

```
VPllexcli:/> security create-host-certificate --host-cert-outfilename TestHostCert.pem  
--host-key-outfilename TestHostKey.pem --host-subject-filenam TestSubject.txt
```

The following files are created in the specified directories:

- /etc/ipsec.d/cacerts/TestHostCert.pem
- /etc/ipsec.d/private/TestHostKey.pem

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-certificate-subject on page 421](#)
- ◆ [security delete-ca-certificate on page 425](#)
- ◆ [security delete-host-certificate on page 426](#)
- ◆ [security export-ca-certificate on page 427](#)
- ◆ [security export-host-certificate on page 428](#)
- ◆ [security import-ca-certificate on page 429](#)
- ◆ [security import-host-certificate on page 431](#)
- ◆ [security ipsec-configure on page 433](#)
- ◆ [security show-cert-subj on page 442](#)
- ◆ [Using the security commands on page 417](#)
- ◆ [EMC VPLEX Security Configuration Guide](#)

## security delete-ca-certificate

Deletes the specified CA certificate and its key.

**Contexts** All contexts.

**Syntax** `security delete-ca-certificate  
[-o|--ca-cert-outfilename] filename  
[-f|--ca-key-outfilename] filename`

**Arguments** **Required arguments**

None (to delete default CA certificate).

**Optional arguments**

`[-o|--ca-cert-outfilename] filename` - CA Certificate output filename.  
Default: strongswanCert.pem.

`[-f|--ca-key-outfilename] filename` - CA Key output filename.

Default: strongswanKey.pem.

**Description** Deletes the CA certificate and deletes the entries from the lockbox that were created by EZ-setup.

**Example** Delete the default CA certificate and key:

```
Vplexcli: /> security delete-ca-certificate
```

The following files are deleted from the specified directories:

- ◆ /etc/ipsec.d/cacerts/strongswanCert.pem
- ◆ /etc/ipsec.d/private/strongswanKey.pem

**Example** Delete a custom CA certificate (not the default):

```
Vplexcli: /> security delete-ca-certificate -o TestCACert.pem -f TestCAKey.pem
```

The following files are deleted from the specified directories:

- ◆ /etc/ipsec.d/cacerts/TestCACert.pem
- ◆ /etc/ipsec.d/private/TestCAKey.pem

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security delete-host-certificate on page 426](#)

## security delete-host-certificate

Deletes the specified host certificate.

**Contexts** All contexts.

**Syntax**  
`security delete-hostcertificate`  
`[-o|--host-cert-outfilename] filename`  
`[-f|--host-key-outfilename] filename`

**Arguments** **Required arguments**  
None (to delete default host certificate).

**Optional arguments**  
`[-o|--host-cert-outfilename] filename` - host certificate output filename.  
Default: hostCert.pem.  
`[-f|--host-key-outfilename] filename` - Host key output filename.  
Default: hostKey.pem.

**Description** Deletes the specified host certificate and deletes the entries from the lockbox that were created by EZ-setup.

**Example** Delete the default host certificate and key:

```
VPLexcli:/> security delete-host-certificate
```

The following files are deleted from the specified directories:

- ◆ /etc/ipsec.d/certs/hostCert.pem
- ◆ /etc/ipsec.d/private/hostKey.pem

**Example** Delete a custom (not the default) host certificate:

```
VPLexcli:/>security delete-host-certificate -o TestHostCert.pem -f TestHostKey.pem
```

The following files are deleted from the specified directories:

- ◆ /etc/ipsec.d/certs/TestHostCert.pem
- ◆ /etc/ipsec.d/private/TestHostKey.pem

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-host-certificate on page 423](#)
- ◆ [security delete-ca-certificate on page 425](#)

---

## security export-ca-certificate

Exports a CA certificate and CA key to a given location.

**Contexts** All contexts.

**Syntax** `security export-ca-certificate`  
`[-c|--ca-cert-filepath] absolute path`  
`[-k|--ca-key-filepath] absolute path`  
`[-e|--ca-export-location] export location`

**Arguments** **Required arguments**

`[-e|--ca-export-location] absolute path` - The absolute path of the location to which to export the CA Certificate and CA Key.

**Optional arguments**

`[-c|--ca-cert-filepath] absolute filepath` - The absolute path of the CA certificate file to export.

Default: `/etc/ipsec.d/cacerts/strongswanCert.pem`.

`[-k|--ca-key-filepath] absolute filepath` - The absolute path of the CA Key file to export.

Default: `/etc/ipsec.d/private/strongswanKey.pem`.

**Description** Exports the CA certificate to the specified location.

---

**Note:** The user executing this command must have write privileges at the location to which the certificate is exported.

---

The import/export of CA certificates does not work for external CA certificates.

**Example** Export the default CA certificate and key to `/var/log/VPlex/cli`:

```
VPlexcli:/> security export-ca-certificate -e /var/log/VPlex/cli
```

**Example** Export a custom CA certificate and its key (created using the “`security create-ca-certificate`” command) to `/var/log/VPlex/cli`:

```
VPlexcli:/> security export-ca-certificate -c /etc/ipsec.d/cacerts/TestCACert.pem -k /etc/ipsec.d/private/TestCAKey.pem -e /var/log/VPlex/cli
```

- See also**
- ◆ [security create-ca-certificate on page 417](#)
  - ◆ [security create-certificate-subject on page 421](#)
  - ◆ [security create-host-certificate on page 423](#)
  - ◆ [security export-host-certificate on page 428](#)
  - ◆ [security import-ca-certificate on page 429](#)
  - ◆ [security import-host-certificate on page 431](#)
  - ◆ [security ipsec-configure on page 433](#)
  - ◆ [security show-cert-subj on page 442](#)
  - ◆ [Using the security commands on page 417](#)
  - ◆ *EMC VPLEX Security Configuration Guide*

## security export-host-certificate

Exports a host certificate and host key to the specified location.

**Contexts** All contexts.

**Syntax**

```
security export-host-certificate
  [-c|--host-cert-filepath] absolute path
  [-k|--host-key-filepath] absolute path
  [-e|--host-export-location] export location
```

**Arguments** **Required arguments**  
[-e|--host-export-location] *absolute path* - The absolute path of the location to which to export the host certificate and host key.

**Optional arguments**  
[-c|--host-cert-filepath] *absolute filepath* - The absolute path of the host certificate file to export.

Default: /etc/ipsec.d/certs/hostCert.pem

[-k|--host-key-filepath] *absolute filepath* - The absolute path of the host key file to export.

Default: /etc/ipsec.d/private/hostKey.pem

**Description** Exports the host certificate to the specified location.

**Note:** The user executing this command must have write privileges at the location to which the certificate is exported.

**Example** Export the default host certificate and key to /var/log/VPlex/cli:

```
VPlexcli:/> security export-host-certificate -e /var/log/VPlex/cli
```

**Example** Export a custom host certificate and it's key (created using the **security create-host-certificate** command) to /var/log/VPlex/cli:

```
VPlexcli:/> security export-host-certificate -c /etc/ipsec.d/certs/TestHostCert.pem -k /etc/ipsec.d/private/TestHostKey.pem -e /var/log/VPlex/cli
```

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-certificate-subject on page 421](#)
- ◆ [security create-host-certificate on page 423](#)
- ◆ [security export-ca-certificate on page 427](#)
- ◆ [security import-ca-certificate on page 429](#)
- ◆ [security show-cert-subj on page 442](#)

## security import-ca-certificate

Imports a CA certificate and CA key from a given location.

**Contexts** All contexts.

**Syntax**

```
security import-ca-certificate
[-c|--ca-cert-filepath] absolute path
[-k|--ca-key-filepath] absolute path
[-i|--ca-cert-import-location] import location
[-j|--ca-key-import-location] import location
```

**Arguments** **Required arguments**

**[-c|--ca-cert-filepath] *absolute path*** The absolute path of the CA certificate file to import.

**[-k|--ca-key-filepath] *absolute path*** - The absolute path of the CA key file to import.

**Optional arguments**

**[-i|--ca-cert-import-location] *import location*** - The absolute path of the location to which to import the CA certificate.

Default location - **/etc/ipsec.d/cacerts**.

**[-j|--ca-key-import-location] *import location*** - The absolute path of the location to which to import the CA certificate.

Default location - **/etc/ipsec.d/private**.

**Description** Imports the CA certificate from the specified location.

**Note:** The user executing this command must have write privileges at the location from which the certificate is imported.

If the import locations for the CA certificate and CA key has files with the same names, they are overwritten.

The import/export of CA certificates does not work for external CA certificates.

**Example** Import the CA certificate and its key from a specified location to the default CA certificate and key location (/var/log/VPlex/cli):

```
VPlexcli:/> security import-ca-certificate -c /var/log/VPlex/cli/strongswanCert.pem -k /var/log/VPlex/cli/strongswanKey.pem
```

- ◆ The imported CA certificate file is copied to /etc/ipsec.d/cacerts
- ◆ The imported CA key is copied to /etc/ipsec.d/private

**Example** Import the CA certificate and key from /var/log/VPlex/cli directory to a custom location:

```
VPlexcli:/> security import-ca-certificate -c /var/log/VPlex/cli/strongswanCert.pem -k /var/log/VPlex/cli/strongswanKey.pem -i /Test/cacerts -j /Test/private
```

- ◆ The imported CA certificate file is copied to /Test/cacerts
- ◆ The imported CA key is copied to /Test/private

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-certificate-subject on page 421](#)
- ◆ [security create-host-certificate on page 423](#)

- ◆ [security export-ca-certificate](#) on page 427
- ◆ [security import-host-certificate](#) on page 431
- ◆ [security ipsec-configure](#) on page 433
- ◆ [security show-cert-subj](#) on page 442

## security import-host-certificate

Imports a host certificate and host key from a given location.

**Contexts** All contexts.

**Syntax**

```
security import-host-certificate
[-c|--host-cert-filepath] absolute path
[-k|--host-key-filepath] absolute path
[-i|--host-cert-import-location] import location
[-j|--host-key-import-location] import location
```

**Arguments** **Required arguments**

**[-c|--host-cert-filepath] *absolute path*** - The absolute path of the host certificate file to import.

**[-k|--host-key-filepath] *absolute path*** - The absolute path of the host key file to import.

**Optional arguments**

**[-i|--host-cert-import-location] *import location*** - The absolute path of the location to which to import the host certificate.

Default location - /etc/ipsec.d/certs.

**[-j|--host-key-import-location] *import location*** - The absolute path of the location to which to import the host certificate.

Default location - /etc/ipsec.d/private.

**Description** Imports the host certificate from the specified location.

**Note:** The user executing this command must have write privileges at the location from which the certificate is imported.

If the import locations for the host certificate and host key have files with the same names, the files are overwritten.

**Example** Import the host certificate and key from /var/log/VPlex/cli:

```
VPlexcli:/> security import-host-certificate -c /var/log/VPlexcli/hostCert.pem -k /var/log/VPlex/cli/hostKey.pem
```

- ◆ The imported host certificate file is copied to /etc/ipsec.d/certs
- ◆ The imported host key is copied to /etc/ipsec.d/private

**Example** Import the host certificate and it's key from /var/log/VPlex/cli to a custom host certificate and key location:

```
VPlexcli:/> security import-ca-certificate -c /var/log/VPlexcli/hostCert.pem -k /var/log/VPlex/cli/hostKey.pem -i /Test/certs -j /Test/private
```

- ◆ The imported host certificate file is copied to /Test/cacerts
- ◆ The imported host key is copied to /Test/private

**See also**

- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-certificate-subject on page 421](#)
- ◆ [security create-host-certificate on page 423](#)
- ◆ [security export-ca-certificate on page 427](#)

- ◆ `security import-ca-certificate` on page 429
- ◆ `security ipsec-configure` on page 433
- ◆ `security show-cert-subj` on page 442

## security ipsec-configure

Configures IPsec after the CA and host certificates have been created.

**Contexts** All contexts.

**Syntax**

```
security ipsec-configure
[-i|--remote-ms-ipaddr] Remote IP address
[-c|--host-cert-filename] host certificate
[-k|--host-key-filename] host key
```

**Arguments** **Required arguments**  
[-i|--remote-ms-ipaddr] *Remote IP address* - IP address of the remote management server.

**Optional arguments**  
[-c|--host-cert-filename] *host certificate* - host certificate filename.  
[-k|--host-key-filename] *host key* - host key filename.

**Description** This command does the following:

- ◆ Backs up the existing ipsec.conf and ipsec.secrets files.
- ◆ Configures ipsec.conf and ipsec.secrets with the latest VPN configuration.
- ◆ Enables the IPsec service at rc3, rc4, and rc5 run levels.
- ◆ Starts the VPN.

The following steps must be completed before using this command:

1. On the first cluster, use the “[security create-ca-certificate](#)” command to create the CA certificate.
2. On the first cluster, use the “[security create-host-certificate](#)” command to create the host certificate.
3. Copy (scp) the CA certificate:  
**From:** /etc/ipsec.d/cacerts/strongswanCert.pem on the first cluster  
**To:** etc/ipsec.d/cacerts/ on the second cluster.
4. Copy (scp) the CA key:  
**From:** /etc/ipsec.d/private/strongswanKey.pem on the first cluster  
**To:** /etc/ipsec.d/cacerts/ on the second cluster.
5. On the second cluster, use the “[security create-host-certificate](#)” command to create the host certificate.
6. On the second cluster, use the “[security ipsec-configure](#)” command, and specify the IP address of the first cluster.
7. On the second cluster, use the “[security ipsec-configure](#)” command. Specify the IP address of the first cluster.
8. On the first cluster, use the “[security ipsec-configure](#)” command. Specify the IP address of the second cluster.
9. On either cluster, use the “[vpn status](#)” command to verify that the VPN is established.

---

**Note:** This command should be used only in VPLEX Metro and VPLEX Geo configurations to create the VPN tunnel between clusters.

---

The distinguished name (DN) used to configure the ipsec is read from the host certificate created on the remote management server. The filename of the host certificate file created on the remote management server must be *hostCert.pem*.

#### About cluster IP seed and cluster ID

The VPLEX IP seed is used to generate the IP addresses used by the internal components of the VPLEX. For more information about components and their IP addresses, refer to *EMC VPLEX Installation and Setup Guide*.

Cluster ID is used by the virtualization software (inter director messaging, cluster identification).

For the current release, the IP seed is the same value as the cluster ID, and always either 1 (cluster-1) or 2 (cluster-2).

**Example** In the following example:

- ◆ On first cluster: the **security ipsec-configure** command configures IPsec to second cluster at 10.6.209.33
- ◆ On second cluster: the **security ipsec-configure** command configures IPsec to first cluster at 10.6.209.32
- ◆ On the first cluster, the **vpn status** command confirms that the VPN to the second cluster is up
- ◆ On the second cluster, the **vpn status** command confirms that the VPN to the first cluster is up

```
Vplexcli-1 Vplexcli:/> security ipsec-configure -i 10.6.209.33
```

```
Vplexcli-2 Vplexcli:/> security ipsec-configure -i 10.6.209.32
```

```
Vplexcli-1 Vplexcli:/> vpn status
```

```
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.6.209.33 is reachable
Remote Internal Gateway addresses are reachable
```

```
Vplexcli-2 Vplexcli:/> vpn status
```

```
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.6.209.32 is reachable
Remote Internal Gateway addresses are reachable
```

- See also**
- ◆ [security create-ca-certificate](#) on page 417
  - ◆ [security create-certificate-subject](#) on page 421
  - ◆ [security create-host-certificate](#) on page 423
  - ◆ [security export-ca-certificate](#) on page 427
  - ◆ [security import-ca-certificate](#) on page 429
  - ◆ [security import-host-certificate](#) on page 431
  - ◆ [security show-cert-subj](#) on page 442
  - ◆ [Using the security commands](#) on page 417
  - ◆ *EMC VPLEX Security Configuration Guide*

---

## security remove-login-banner

Removes the login banner from the management server.

**Contexts** All contexts.

**Syntax** `security remove-login-banner  
[-f|--force]`

**Arguments** **Required arguments**  
None.

**Optional arguments**

`[-f|--force]` - Forces the removal of the login banner without asking for any user confirmation. Allows this command to be run from non-interactive scripts.

**Description** Removes a custom login banner from the management server.  
The change takes effect at the next login to the management server.

**Example** Remove the login banner:

```
Vplexcli: /> security remove-login-banner
```

```
The login banner of this management server will be removed.  
Do you want to proceed ? (Yes/No) yes
```

**See also** ♦ [security set-login-banner on page 441](#)

## security renew-all-certificates

Renews CA and host security certificates.

**Contexts** All contexts.

**Syntax** `security renew-all-certificates`

**Arguments** None.

**Description** When VPLEX is installed, EZ-Setup creates one CA certificate and two or three host certificates:

- ◆ Certification Authority (CA) certificate shared by all clusters
- ◆ VPN host certificate
- ◆ Web server host certificate
- ◆ VPLEX Witness host certificate (when VPLEX Witness is installed)

All types of certificates expire and must be periodically renewed. By default:

- ◆ CA certificates must be renewed every 5 years
- ◆ Host certificates must be renewed every 2 years

Starting in Release 5.0.1 Patch, use the **security renew-all-certificates** CLI command to renew all security certificates on a VPLEX system.

In Metro and Geo systems, run the command twice, once on each cluster. For systems with VPLEX Witness deployed, make sure you run the command first on the cluster where VPLEX Witness was initially installed. See [Before you begin: on page 437](#) for the steps to determine the correct cluster.

You can use the command at any time to renew certificates whether or not they are about to expire.

Each certificate has an associated passphrase. For releases prior to 4.2, VPLEX was not aware of any passphrases. For releases after 4.2, VPLEX may or may not be aware of passphrases. Previous to 5.0.1 Patch, VPLEX is never aware of the passphrase for the VPLEX Witness certificate. During renewal, you are prompted to enter any passphrases that VPLEX does not know.

After renewal, VPLEX is aware of all passphrases.

There are two general methods for renewing passphrases:

- ◆ Renew the security certificates using their **current** passphrases.

VPLEX may find the current passphrases for the CA, VPN, and Web server security certificates, but prior to 5.0.1 Patch, never finds the current passphrase for the VPLEX Witness security certificate.

If you choose to renew the certificates using their current passphrases, you are prompted to provide the passphrase for any certificate that VPLEX does not find.



### **IMPORTANT**

**In Metro and Geo systems, you are always prompted for the Certificate Authority (CA) passphrase when you run the command on the second cluster.**

- ◆ Renew the certificates using a **common** passphrase.

All certificates are renewed using the same passphrase.



## CAUTION

In Metro and Geo systems, do not renew the security certificates using the current passphrases if you do not have a record of the Certificate Authority (CA) passphrase. You must provide the current CA passphrase when you renew the certificates on the second cluster. If you do not have a record of the CA passphrase, do not renew the certificates until you have the passphrase or renew with a common passphrase.

### Example: Renew certificates on a VPLEX Local

The following is the simplest renewal case:

- ◆ VPLEX Local
- ◆ Using a common passphrase

```
Vplexcli:/> security renew-all-certificates
```

If this is a VPLEX Local deployment, follow the prompts below to renew all certificates. If this is a VPLEX Metro or Geo, please contact EMC Support.

```
Do you want to continue renewing all certificates? (Y/N): yes
```

```
Detecting all the VPLEX certificates currently configured on the system...<
```

```
The following certificates will be renewed:
```

| Certificate Type           | Expiration Date          | New Expiration Date      |
|----------------------------|--------------------------|--------------------------|
| Host Certificate (VPN)     | Sep 11 16:22:22 2013 GMT | Sep 11 16:22:22 2015 UTC |
| Certificate Authority (CA) | Sep 10 16:22:21 2016 GMT | Sep 9 16:22:21 2021 UTC  |
| Host Certificate (WEB)     | Sep 11 16:22:24 2013 GMT | Sep 11 16:22:24 2015 UTC |

```
The certificates above will be renewed, to expire on the dates shown. Do you want to continue? (Y/N): y
```

```
Would you like to renew the certificates using the current passphrases? (Y/N): no
```

```
Please create a passphrase (at least 8 chars) to be used for all the certificate renewals:
```

```
Re-enter the passphrase for the Certificate Key: CA-passphrase
```

```
Renewing CA certificate...
```

```
The CA certificate was successfully renewed.
```

```
Renewing VPN certificate...
```

```
The VPN certificate was successfully renewed.
```

```
Renewing WEB certificate...
```

```
Your Java Key Store has been created.
```

```
https keystore: /var/log/Vplex/cli/.keystore
```

```
started web server on ports {'http': 49880, 'https': 49881}
```

```
The Web certificate was successfully renewed.
```

```
Generating certificate renewal summary...
```

```
All VPLEX certificates have been renewed successfully
```

### Example: Renew certificates on a VPLEX Geo with VPLEX Witness

The following example renews security certificates in a more complex scenario:

- VPLEX Geo
- Using current passphrases
- VPLEX does not find any current passphrases
- VPLEX Witness is installed

**Before you begin:**

- ◆ Navigate to the `/ect/ssl` directory on the management servers, and see on which cluster the file `index.txt` includes this string: `CN=Vplex VPN CWS`. If the string is present, run the renewal command on that cluster first.

- ◆ Use the **vpn status** command to verify that the VPN tunnel between clusters is operational, and the Cluster Witness Server is reachable. Do not proceed if these conditions are not present.
- ◆ Use the **ll cluster-witness** command to verify that the cluster-witness admin-state is disabled. If it is enabled, use the **cluster-witness disable** command to disable it.

### On the cluster where the VPLEX Witness certificate was created

```
VPLexcli:/> security renew-all-certificates
```

Please note that to renew certificates on a Metro or Geo deployment, this command must be run on both clusters.

WARNING : After running this command on the first cluster, the VPN tunnel between clusters will be down temporarily until you run this command on the second cluster. This will not affect I/O but will result in the inability to manage the remote cluster.

Detecting all the VPLEX certificates currently configured on the system...

The following certificates will be renewed:

| Certificate Type           | Expiration Date          | New Expiration Date      |
|----------------------------|--------------------------|--------------------------|
| Host Certificate (VPN)     | Sep 12 17:10:10 2013 GMT | Sep 12 17:10:10 2015 UTC |
| Host Certificate (CW)      | Sep 12 17:10:18 2013 GMT | Sep 12 17:10:18 2015 UTC |
| Certificate Authority (CA) | Sep 11 17:10:08 2016 GMT | Sep 10 17:10:08 2021 UTC |
| Host Certificate (WEB)     | Sep 12 17:10:16 2013 GMT | Sep 12 17:10:16 2015 UTC |

The certificates above will be renewed, to expire on the dates shown. Do you want to continue?(Y/N): **y**

Would you like to renew the certificates using the current passphrases? (Y/N): **y**

Some or all of the passphrases are not available, so new passphrases must be created:

Please create a passphrase (at least 8 chars) for the Certificate Authority renewal:

**CA-passphrase**

Re-enter the passphrase for the Certificate Key: **CA-passphrase**

Please create a passphrase (at least 8 chars) for the VPN certificate renewal: **VPN-passphrase**

Re-enter the passphrase for the Certificate Key: **VPN-passphrase**

Please create a passphrase (at least 8 chars) for the web certificate renewal: **WEB-passphrase**

Re-enter the passphrase for the Certificate Key: **WEB-passphrase**

Please create a passphrase (at least 8 chars) for the cluster witness certificate renewal:

**CWS-passphrase**

Re-enter the passphrase for the Certificate Key: **CWS-passphrase**

Renewing CA certificate...

The CA certificate was successfully renewed.

Renewing VPN certificate...

The VPN certificate was successfully renewed.

Renewing WEB certificate...

Your Java Key Store has been created.

https keystore: /var/log/VPLEX/cli/.keystore

started web server on ports {'http': 49880, 'https': 49881}

The Web certificate was successfully renewed.

Renewing CW certificate...  
The CWS certificate was successfully renewed.

Generating certificate renewal summary...  
Certificates have been successfully renewed on this cluster. To complete the renewal process, run this command on the second cluster.

\* The VPN tunnel between clusters will be down temporarily

### On the second cluster

VPlexcli:/> **security renew-all-certificates**

Please note that to renew certificates on a Metro or Geo deployment, this command must be run on both clusters.

WARNING : After running this command on the first cluster, the VPN tunnel between clusters will be down temporarily until you run this command on the second cluster. This will not affect I/O but will result in the inability to manage the remote cluster.

Before continuing to renew certificates on this cluster, please confirm that certificates have been renewed on the other cluster.

Have certificates have been renewed on the other cluster? (yes/no) (Y/N): **y**  
Detecting all the VPLEX certificates currently configured on the system...

The following certificates will be renewed:

| Certificate Type           | Expiration Date          | New Expiration Date      |
|----------------------------|--------------------------|--------------------------|
| Host Certificate (VPN)     | Sep 12 17:13:04 2013 GMT | Sep 12 17:13:04 2015 UTC |
| Certificate Authority (CA) | Sep 11 17:10:08 2016 GMT | Sep 10 17:10:08 2021 UTC |
| Host Certificate (WEB)     | Sep 12 17:13:09 2013 GMT | Sep 12 17:13:09 2015 UTC |

The certificates above will be renewed, to expire on the dates shown. Do you want to continue? (Y/N): **y**

Would you like to renew the certificates using the current passphrases? (Y/N): **y**

Some or all of the passphrases are not available, so new passphrases must be created:

Please enter the 'service' account password( 8 chars ) for the Remote Management Server:  
**Mi@Dim7T**

Re-enter the password: **Mi@Dim7T**

Please enter the passphrase for the Certificate Authority on the remote cluster: **CA-passphrase**

Re-enter the passphrase for the Certificate Key: **CA-passphrase**

Please create a passphrase (at least 8 chars) for the VPN certificate renewal: **VPN-passphrase**

Re-enter the passphrase for the Certificate Key: **VPN-passphrase**

Please create a passphrase (at least 8 chars) for the web certificate renewal: **WEB-passphrase**

Re-enter the passphrase for the Certificate Key: **WEB-passphrase**

Renewing CA certificate...  
The CA certificate was successfully renewed.

Renewing VPN certificate...  
The VPN certificate was successfully renewed.

Renewing WEB certificate...

```
Your Java Key Store has been created.
https keystore: /var/log/VPlex/cli/.keystore
started web server on ports {'http': 49880, 'https': 49881}
The Web certificate was successfully renewed.
```

Generating certificate renewal summary...

### **After certificate renewal: enable VPLEX Witness**

If VPLEX Witness was disabled before the security certificates were renewed:

- ◆ Use the **cluster-witness enable** command to re-enable VPLEX Witness.
- ◆ Use the **ll cluster-witness** command to verify that the admin-state is enabled.

#### **See also**

- ◆ [cluster-witness disable on page 98](#)
- ◆ [cluster-witness enable on page 101](#)
- ◆ [security create-ca-certificate on page 417](#)
- ◆ [security create-host-certificate on page 423](#)
- ◆ [security export-ca-certificate on page 427](#)
- ◆ [security export-host-certificate on page 428](#)
- ◆ [security import-ca-certificate on page 429](#)
- ◆ [security import-host-certificate on page 431](#)

## security set-login-banner

Applies a text file as the login banner on the management server.

**Contexts** All contexts.

**Syntax**  
`security set-login-banner  
[-b|--login-banner-file] file  
[-f|--force]`

**Arguments** **Required arguments**

**[-b|--login-banner-file] file** - Full pathname to the file containing the formatted login banner text.

**Optional arguments**

**[-f|--force]** - Forces the addition of the login banner without asking for any user confirmation. Allows this command to be run from non-interactive scripts.

**Description** This command sets the login banner for the management server. This command applies the contents of the specified text file as the login banner.

The change takes effect at the next login to the management server.

The formatting of the text in the specified text file is replicated in the banner.

There is no limit to the number of characters or lines in the specified text file.

Use this command to create a customized login banner. The formatting of the text in the specified text file is replicated in the banner.

**Example** In the following example, a text file “login-banner.txt” containing the following lines is specified as the login banner:

```
VPLEX cluster-1/Hopkinton
```

```
Test lab 3, Room 6, Rack 47
```

```
Metro with RecoverPoint CDP
```

```
Vplexcli:/> security set-login-banner -b  
/home/service/login-banner.txt
```

```
The text provided in the specified file will be set as the Login banner  
for this management server.
```

```
Any previously applied banner will be overwritten.
```

```
Do you want to proceed ? (Yes/No) Yes
```

At next login to the management server, the new login banner is displayed:

```
login as: service  
VPLEX cluster-1/Hopkinton  
Test lab 3, Room 6, Rack 47  
Metro with RecoverPoint CDP  
Password:
```

**See also** ♦ [security remove-login-banner on page 435](#)

## security show-cert-subj

Displays the certificate subject file.

**Contexts** All contexts.

**Syntax** `security show-cert-subj  
[s|--subject-infilename] filename`

**Arguments** **Required arguments**  
[**-s** | **--subject-infilename**] *filename* - Filename of the certificate subject file to display.  
The file is assumed to reside in the following directory on the management server:  
`/var/log/VPlex/cli`

**Description** Displays the certificate subject file.

**Example**

```
VPlexcli: /> security show-cert-subj -s CACertSubjectInfo.txt
SUBJECT_COUNTRY=US
SUBJECT_STATE=Massachusetts
SUBJECT_LOCALITY=Hopkinton
SUBJECT_ORG=EMC
SUBJECT_ORG_UNIT=EMC
SUBJECT_COMMON_NAME=FN00094400134
SUBJECT_EMAIL=support@emc.com
```

**See also** ♦ [security create-certificate-subject on page 421](#)

---

## sessions

Displays active Unisphere for VPLEX sessions.

**Contexts** All contexts.

**Syntax** sessions

**Description** Displays the username, hostname, port and start time of active sessions to the Unisphere for VPLEX.

**Example**

```
Vplexcli:/> sessions
```

| Type          | Username | Hostname  | Port  | Creation Time                |
|---------------|----------|-----------|-------|------------------------------|
| TELNET_SHELL  | service  | localhost | 23848 | Wed Sep 15 15:34:33 UTC 2010 |
| DEFAULT_SHELL | -        | -         | -     | Tue Aug 03 17:16:07 UTC 2010 |

## set

Changes the value of a writable attribute(s) in the given context.

**Contexts** All contexts.

**Syntax**

```
set
  [-d|--default]
  [-f|--force]
  [-a|--attributes] selector pattern
  [-v|--value] new value
```

### Optional arguments

**[-d|--default]** - Sets the specified attribute(s) to the default value(s), if any exist. If no attributes are specified, displays the default values for attributes in the current/specified given context.

**[-f|--force]** - Force the value to be set, bypassing any confirmations or guards.

**[-a|--attributes]** *selector pattern* - \* Attribute selector pattern.

**[-v|--value]** *new value* - \* The new value to assign to the specified attribute(s).

\* - argument is positional.

**Description** Use the **set** command with no arguments to display the attributes available in the current context.

Use the **set --default** command with no additional arguments to display the default values for the current context or a specified context.

Use the **set** command with an attribute pattern to display the matching attributes and the required syntax for their values.

Use the **set** command with an attribute pattern and a value to change the value of each matching attribute to the given value.

An attribute pattern is an attribute name optionally preceded with a context glob pattern and a double-colon (:). The pattern matches the named attribute on each context matched by the glob pattern.

If the glob pattern is omitted, **set** assumes the current context.

If the value and the attribute name are omitted, **set** displays information on all the attributes on all the matching contexts.

**Example** Display which attributes are writable in the current context, and their valid inputs:

```
Vplexcli:/distributed-storage/distributed-devices/TestDisDevice> set
attribute          input-description
-----
application-consistent Takes one of '0', '1', 'f', 'false', 'n', 'no', 'off', 'on', 't', 'true', 'y',
'yes' (not case sensitive).
auto-resume         Takes one of '0', '1', 'f', 'false', 'n', 'no', 'off', 'on', 't', 'true', 'y',
'yes' (not case sensitive).
block-count          Read-only.
block-size            Read-only.
capacity              Read-only.
clusters-involved     Read-only.
.
.
.
```

Use the **--default** argument without any attribute(s) to display the default values for the current (or specified) context's attributes:

```
Vplexcli:/distributed-storage/distributed-devices/TestDisDevice> set --default
attribute          default-value
```

```

-----
application-consistent  No default value.
auto-resume             No default value.
block-count             No default value.
.
.
.

```

**Example** Change the name of a meta-volume:

```

VPlexcli:/clusters/cluster-1/system-volumes/new_meta1_backup_2010May24_163810> set name backup_May24_pre_refresh

```

**Example** Display information about attributes in the eth0 context:

```

VPlexcli:/> set /management-server/ports/eth0
attribute                                     input-description
-----
/mangement-server/ports/eth0::address        Read-only.
/mangement-server/ports/eth0::gateway        Read-only.
/mangement-server/ports/eth0::name           Read-only.
/mangement-server/ports/eth0::net-mask       Read-only

```

**Example** Display the address attribute for eth0:

```

VPlexcli:/> set /management-server/ports/eth0::address
attribute                                     input-description
-----
/mangement-server/ports/eth0::address        Read-only.

```

**Example** Set the remote IP address and started attributes for SNMP traps:

```

VPlexcli:/notifications/call-home/snmp-traps/Test> set remote-host 10.6.213.39
VPlexcli:/notifications/call-home/snmp-traps/Test> set started true

```

**Example** Attach a rule-set 'cluster1\_Active' to the 'device dd\_00':

```

VPlexcli:/distributed-storage/distributed-devices> set dd_00::rule-set-name cluster1_Active

```

**Example** Set a storage volume's thin-rebuild attribute to true:

```

VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes/clar_LUN83> set thin-rebuild true
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes/clar_LUN83> ll
Name                                     Value
-----
application-consistent  false
.
.
.
storage-volumetype       normal
system-id                 VPD83T3:6006016091c50e005057534d0c17e011
thin-rebuild            true
total-free-space         2G
use                       claimed
used-by                   []
vendor-specific-name      DGC

```

**Enable/disable call-home notifications** Use the **set enabled false --force** command in notifications/call-home context to disable call-home notifications (recommended during disruptive operations):

```

VPlexcli:/> cd /notifications/call-home/
VPlexcli:/notifications/call-home> set enabled false --force

```

Use the **set enabled true** command in notifications/call-home context to enable call-home notifications:

```
VPLEXcli: /> cd /notifications/call-home/
```

```
VPLEXcli: /notifications/call-home> set enabled true
```

- See also**
- ◆ [notifications call-home test on page 365](#)
  - ◆ [set topology on page 447](#)
  - ◆ [storage-volume claim on page 458](#)
  - ◆ [storage-volume unclaim on page 474](#)

## set topology

Changes the topology attribute for a Fibre Channel port.

**Contexts** /engines/engine/directors/director/hardware/ports/port

**Syntax** set topology [p2p|loop]

**Arguments** **Required arguments**

**p2p** - Sets the port's topology as point-to-point. The port comes up as an F-port.

Use the p2p topology to connect the Fibre Channel fabric to a node.

**loop** - Sets the port's topology as loop. The port comes up as an FL-Port.

Use the loop topology to connect a Fibre Channel Arbitrated Loop (ring-style network topology) to a fabric.

**Description** Change the default setting for a Fibre Channel port.

Default: p2p.

**Example** Navigate to a FC port context and set the topology as p2p:

```
VPLEXcli: /> cd /engines/engine-1-1/directors/Cluster_1_Dir1A/hardware/ports/A4-FC02
VPLEXcli:/engines/engine-1-1/directors/Cluster_1_Dir1A/hardware/ports/A4-FC02> set topology
p2p
```

```
VPLEXcli:/engines/engine-1-1/directors/Cluster_1_Dir1A/hardware/ports/A4-FC02> ll
```

| Name               | Value              |
|--------------------|--------------------|
| address            | 0x5000144240014742 |
| current-speed      | 8Gbits/s           |
| description        | -                  |
| enabled            | true               |
| max-speed          | 8Gbits/s           |
| node-wwn           | 0x500014403ca00147 |
| operational-status | ok                 |
| port-status        | up                 |
| port-wwn           | 0x5000144240014742 |
| protocols          | [fc]               |
| role               | wan-com            |
| target-port        | -                  |
| topology           | p2p                |

**See also** ♦ [set on page 444](#)

## sms dump

Collects the logs files on the management server.

**Contexts** All contexts.

**Syntax**  
sms dump  
[-d|--destination-directory] *destination directory*  
[-t|--target\_log] *target logName*

**Arguments** **Required arguments**  
[-d|--destination-directory] *destination directory* - Destination directory for the sms dump logs.

**Optional arguments**  
[-t|--target\_log] *target logName* - Collect only files specified under logName from smsDump.xml.

**Description** Collects the following log files:

### Cliilogs

/var/log/VPlex/cli/client.log\* -- VS1CLI logs, logs dumped by VS1cli scripts

/var/log/VPlex/cli/session.log\* -- what the user does in a VS1CLI session

/var/log/VPlex/cli/firmware.log\* -- nsfw.log files from all directors

### Connectemc

/var/log/ConnectEMC/logs/\* -- connectemc logs

/opt/emc/connectemc/archive -- connectemc logs

/opt/emc/connectemc/failed -- connectemc logs

/opt/emc/connectemc/\*.xml -- connectemc logs

/opt/emc/connectemc/\*.ini -- connectemc logs

/var/log/VPlex/cli/ema\_adaptor.log\*

### Configuration

/var/log/VPlex/cli/\*.config

/var/log/VPlex/cli/\*.xml

/var/log/VPlex/cli/\*.properties

/var/log/VS1/cli/persistentstore.xml -- generated when user connects to VS1CLI

/var/log/VPlex/cli/connections -- what the VS1CLI is connected to.

/var/log/VPlex/cli/VPlexcommands.txt

/var/log/VPlex/cli/VPlexconfig.xml

/var/log/VPlex/cli/VPlexcli-init

/opt/vs1/backup/\*.ini

/opt/vs1/backup/\*.xml

/opt/emc/VPlex/\*.xml

/opt/emc/VPlex/\*.properties

### Upgrade

/var/log/VPlex/cli/capture/\* (ndu status files)  
/tmp/VPlexInstallPackages/\*.xml  
/tmp/VPlexInstallPackages/\*.properties  
/tmp/VPlexInstallPackages/\*.log  
/var/log/install.log

### System

/var/log/warn\*  
/var/log/messages\*  
/var/log/boot.msg  
/var/log/boot.omsg  
/var/log/firewall  
/etc/sysconfig/SuSEfirewall2  
/etc/sysconfig/network/ifcfg\*  
/etc/sysconfig/network/ifroute\*  
/etc/sysctl.conf

**Example** Collect the logs files on the management server and send them to the designated directory:

```
VPlexcli:> sms dump --destination-directory /var/log/VPlex/cli  
Initiating sms dump...  
sms dump completed to file  
  /var/log/VPlex/cli/smsDump_2010-09-15_16.40.20.zip.
```

**See also**

- ◆ [cluster configdump on page 81](#)
- ◆ [collect-diagnostics on page 106](#)
- ◆ [director appdump on page 190](#)
- ◆ [getsysinfo on page 272](#)

## snmp-agent configure

Configures the VPLEX SNMP agent service on the local cluster.

**Contexts** All contexts.

**Syntax** `snmp-agent configure`

**Description** Configures SNMP agent on the local cluster, and starts the SNMP agent. **snmp-agent configure** checks the number of directors in the local cluster and configures the VPLEX SNMP agent on the VPLEX management server. Statistics can be retrieved from all directors in the local cluster.

**Note:** All the directors have to be operational and reachable through the VPLEX management server before the SNMP agent is configured. When configuration is complete, the VPLEX `snmp-agent` starts automatically.

The VPLEX SNMP agent:

- ◆ Supports retrieval of performance-related statistics as published in the VPLEX-MIB.mib.
- ◆ Runs on the management server and fetches performance related data from individual directors using a firmware specific interface.
- ◆ Provides SNMP MIB data for directors for the local cluster only.
- ◆ Runs on Port 161 of the management server and uses the UDP protocol.
- ◆ Supports the following SNMP commands:
  - SNMP Get
  - SNMP Get Next
  - SNMP get Bulk

The SNMP Set command is not supported in this release.

VPLEX supports SNMP version `snmpv2c`.

VPLEX MIBs are located on the management server in the `/opt/emc/Vplex/mibs` directory.

Use the public IP address of the VPLEX management server to retrieve performance statistics using SNMP.

**Example**

```
Vplexcli:/> snmp-agent configure
The community string is already configured to be: private.
Choosing to continue will change the existing community string.

Do you want to continue? (yes/no)yes

What community string should the agent use? [private]: public

Vplexcli:/>
```

- See also**
- ◆ [snmp-agent start on page 451](#)
  - ◆ [snmp-agent status on page 452](#)
  - ◆ [snmp-agent stop on page 453](#)
  - ◆ [snmp-agent unconfigure on page 454](#)

---

## snmp-agent start

Starts the SNMP agent service.

**Contexts** All contexts.

**Syntax** `snmp-agent start`

**Description** Starts the SNMP agent on the local cluster.

The SNMP agent must be configured before this command can be used.

**Example**  
Vplexcli:/> **snmp-agent start**  
SNMP agent has been started.

**See also**

- ◆ [snmp-agent configure on page 450](#)
- ◆ [snmp-agent status on page 452](#)
- ◆ [snmp-agent stop on page 453](#)
- ◆ [snmp-agent unconfigure on page 454](#)

---

## snmp-agent status

Displays the SNMP agent service on the local cluster.

**Contexts** All contexts.

**Syntax** `snmp-agent status`

**Description** Displays the status of the SNMP agent on the local cluster.

**Example** SNMP agent is running:

```
Vplexcli:/> snmp-agent status  
SNMP Agent Service status is: Running
```

SNMP agent is not configured (not running):

```
Vplexcli:/> snmp-agent status  
SNMP Agent Service status is: Unconfigured
```

- See also**
- ◆ [snmp-agent configure on page 450](#)
  - ◆ [snmp-agent start on page 451](#)
  - ◆ [snmp-agent stop on page 453](#)
  - ◆ [snmp-agent unconfigure on page 454](#)

---

## snmp-agent stop

Stops the SNMP agent service.

**Contexts** All contexts.

**Syntax** `snmp-agent stop`

**Description** Stops the SNMP agent on the local cluster.

The SNMP agent must be configured before this command can be used.

**Example**  
Vplexcli:/> **snmp-agent stop**  
SNMP agent has been stopped.

**See also**

- ◆ [snmp-agent configure on page 450](#)
- ◆ [snmp-agent status on page 452](#)
- ◆ [snmp-agent start on page 451](#)
- ◆ [snmp-agent unconfigure on page 454](#)

---

## snmp-agent unconfigure

Destroys the SNMP agent.

**Contexts** All contexts.

**Syntax** `snmp-agent unconfigure`

**Description** Unconfigures the SNMP agent on the local cluster, and stops the agent.

**Example**  
VPLEXcli: /> **snmp-agent unconfigure**  
SNMP agent has been unconfigured.

- See also**
- ◆ [snmp-agent configure on page 450](#)
  - ◆ [snmp-agent start on page 451](#)
  - ◆ [snmp-agent status on page 452](#)
  - ◆ [snmp-agent stop on page 453](#)

## source

Reads and executes commands from a script.

**Contexts** All contexts.

**Syntax** `source`  
`[-f|--file] filename`

**Arguments** **Required arguments**  
`[-f|--file] filename` - \* Name of the script file to read and execute.  
\* - argument is positional.

**Description** Filenames use the syntax of the underlying platform.  
The script file may contain any CLI commands.  
If the **exit** command is included, the shell exits immediately, without processing the commands that follow it in the file.

**Example** In the following example, a text file Source.txt contains only two commands:

```
service@ManagementServer:/var/log/VPlex/cli> cat Source.txt
version -a
exit
When executed:
The first command in the file is run
The exit command exits the command shell
VPlexcli:/> source --file /var/log/VPlex/cli/Source.txt
What                                     Version                                Info
-----
Product Version                         4.1.0.00.00.12                        -
SMSv2                                    0.16.15.0.0                            -
Mgmt Server Base                         D4_MSB_7                                -
Mgmt Server Software                     D4.70.0.9                               -
/engines/engine-2-1/directors/Cluster_2_Dir_1B 1.2.43.9.0                             -
/engines/engine-2-1/directors/Cluster_2_Dir_1A 1.2.43.9.0                             -
/engines/engine-1-1/directors/Cluster_1_Dir1B 1.2.43.9.0                             -
/engines/engine-1-1/directors/Cluster_1_Dir1A 1.2.43.9.0                             -
/engines/engine-2-2/directors/Cluster_2_Dir_2B 1.2.43.9.0                             -
/engines/engine-2-2/directors/Cluster_2_Dir_2A 1.2.43.9.0                             -

Connection closed by foreign host.
service@ManagementServer:~>
```

**See also** ♦ [script on page 416](#)

## storage-volume auto-unbanish-interval

Displays or changes auto-unbanish interval on a single director.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **auto-unbanish-interval**.

**Syntax**

```
storage-volume auto-unbanish-interval
[-n|--director] context path
[-i|--interval] [20 seconds and greater]
```

**Arguments** **Required arguments**

`[-n|--director] context path` - \* The director on which to show or change the delay for automatic unbanishment.

**Optional arguments**

`[-i|--interval] [20 seconds and greater]` - Number of seconds the director firmware waits before unbanishing a banished storage-volume (LUN).

**Range:** 20 seconds - no upper limit.

**Default:** 30 seconds.

\* - argument is positional.

**Description** See [Banished storage volumes \(LUNs\) on page 472](#).

At regular intervals, the VPLEX directors looks for logical units that were previously banished. If VPLEX finds banished logical units, it unbanishes them. This process happens automatically and continuously, and includes a delay interval with a default value of 30 seconds.

Every 30 seconds the process looks for previously banished logical units and unbanishes any it finds.

Use this command to display and/or change the delay interval.

---

**Note:** This change in the interval value is not saved between restarts of the director firmware (NDU, director reboots). When the director firmware is restarted, the interval value is reset to the default of 30 seconds.

---

Use the **auto-unbanish-interval --director *director*** command to display the current delay (in seconds) for automatic unbanishment on the specified director.

Use the **auto-unbanish-interval --director *director* --interval *interval*** command to change the delay timer for the specified director to the specified number of seconds.

The default metric for setting the **--interval** argument is seconds, but minutes and hours, and days are accepted. The following are valid values for the **--interval** argument: 2s, 2second, 2seconds, 2sec, 2min, 2minute, 2minutes, 2hr, 2hours, 2hour.

---

**Note:** The interval is displayed in seconds.

---

**Example** In the following example:

- ◆ The **auto-unbanish-interval --director *director* --interval *interval*** command changes the delay timer to 200 seconds.

- ◆ The `auto-unbanish-interval --director director` command displays the new setting.

```
Vplexcli:/> storage-volume auto-unbanish-interval --director director-1-1-A --interval 200
```

```
Vplexcli:/> storage-volume auto-unbanish-interval --director director-1-1-A  
200 seconds
```

- See also**
- ◆ [storage-volume list-banished on page 466](#)
  - ◆ [storage-volume unbanish on page 472](#)

## storage-volume claim

Claims the specified storage volumes.

**Contexts** All contexts.

In /clusters/*cluster*/storage-elements/storage-volumes context, command is **claim**.

**Syntax**

```
storage-volume claim
  [--appc]
  [-n|--name] new name
  --thin-rebuild
  --batch-size integer
  [-d|--storage-volumes] context path,context path...
```

**Arguments** **Required arguments**

**[-d|--storage-volumes] *context path,context path...*** - \* List of one or more storage volumes to claim.

**Optional arguments**

**[--appc]** - Make the specified storage volumes application consistent. Prevents data already on the specified storage volume(s) from being deleted or overwritten during the process of constructing a virtual volume.

After a virtual volume is constructed using this storage volume, there is no restriction on the access to the data, i.e. the data can be overwritten by host I/O.



### CAUTION

**The application consistent attribute may be modified using the “set” command but only when the storage-volume is in the claimed state. The application consistent attribute may not be altered for storage volumes that are unclaimed or in use.**

**[-n|--name] *new name*** - The new name of the storage-volume after it is claimed.

**--thin-rebuild** - Claims the specified storage volumes as “thin”. Thin storage allocates blocks of data on demand versus allocating all the blocks up front.

If a storage volume has already been claimed, it can be designated as “thin” using the “set” command.

**--batch-size *integer*** - When using wildcards to claim multiple volumes with one command, the maximum number of storage volumes to claim at once.

\* - argument is positional.

**Description**

A storage volume is a device or LUN that is visible to VPLEX. The capacity of storage volumes is used to create extents, devices and virtual volumes.

Storage volumes must be claimed, and optionally named before they can be used in a VPLEX cluster. Once claimed, the storage volume can be used as a single extent occupying the volume’s entire capacity, or divided into multiple extents (up to 128).

### Thin provisioning

Thin provisioning allows storage to migrate onto a thinly provisioned storage volumes while allocating the minimal amount of thin storage pool capacity.

Thinly provisioned storage volumes can be incorporated into RAID 1 mirrors with similar consumption of thin storage pool capacity.

VPLEX preserves the unallocated thin pool space of the target storage volume by detecting zeroed data content before writing, and suppressing the write for cases where it would cause an unnecessary allocation. VPLEX requires the user to specify thin provisioning for each back-end storage volume. If a storage volume is thinly provisioned, the "thin-rebuild" attribute must be to "true" either during or after claiming.



### CAUTION

**If a thinly provisioned storage volume contains non-zero data before being connected to VPLEX, the performance of the migration or initial RAID 1 rebuild is adversely affected**

**System volumes are supported on thinly provisioned LUNs, but these volumes must have their full capacity of thin storage pool resources set aside and not be in competition for this space with any user-data volumes on the same pool.**

If:

- ◆ The thin storage allocation pool runs out of space, and
- ◆ If this is the last redundant leg of the RAID 1,

further writing to a thinly provisioned device causes the volume to lose access to the device, a DU.

**Example** In the following example:

- ◆ The **ll** command in storage-volumes context displays the available storage.
- ◆ The **claim** command claims the specified unclaimed storage volume from the clusters/*cluster*/storage-elements/storage-volumes context.

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes>ll
```

```
.
```

| Name                        | VPD83 ID                                 | Capacity | Use     | Vendor | IO     | Type   | Thin    |
|-----------------------------|------------------------------------------|----------|---------|--------|--------|--------|---------|
| -----                       | -----                                    | -----    | -----   | -----  | Status | -----  | Rebuild |
| -----                       | -----                                    | -----    | -----   | -----  | -----  | -----  | -----   |
| Basic_c1_ramdisk_100GB_684_ | VPD83T3:60001440000000103017dfea88355431 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_685_ | VPD83T3:60001440000000103017dfea88355433 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_686_ | VPD83T3:60001440000000103017dfea88355435 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_687_ | VPD83T3:60001440000000103017dfea88355437 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_688_ | VPD83T3:60001440000000103017dfea88355439 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_689_ | VPD83T3:60001440000000103017dfea8835543b | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_68_  | VPD83T3:60001440000000103017dfea88354f61 | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_690_ | VPD83T3:60001440000000103017dfea8835543d | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_691_ | VPD83T3:60001440000000103017dfea8835543f | 100G     | claimed | EMC    | alive  | normal | false   |
| Basic_c1_ramdisk_100GB_692_ | VPD83T3:60001440000000103017dfea88355441 | 100G     | claimed | EMC    | alive  | normal | false   |

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes>claim --storage-volumes
VPD83T3:6006016021d025007029e95b2327df11
```

**Example** Claim a storage volume and name it Symm1254\_7BF from the clusters/*cluster* context:

```
Vplexcli:/clusters/cluster-1> storage-volume claim -name Symm1254_7BF -d
VPD83T3:60000970000192601254533030374241
```

**Example** Claim storage volumes using the **--thin-rebuild** option. In the following example:

- ◆ The **claim** command with **--thin-rebuild** claims two storage volumes as thin storage (from the clusters/*cluster*/storage-elements/storage-volumes context)
- ◆ The **ll** command displays one of the claimed storage volumes:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> claim --thin-rebuild --storage-volumes
```

```
VPD83T3:6006016091c50e005057534d0c17e011,VPD83T3:6006016091c50e005257534d0c17e011
```

Of the 2 storage-volumes that were given, 2 storage-volumes were claimed.

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> ll
```

```
VPD83T3:6006016091c50e005057534d0c17e011
```

```
/clusters/cluster-1/storage-elements/storage-volumes/VPD83T3:6006016091c50e005057534d0c17e011  
:
```

| Name                   | Value                                                                                                                                                                                                                                                                                                                                                       |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| application-consistent | false                                                                                                                                                                                                                                                                                                                                                       |
| block-count            | 524288                                                                                                                                                                                                                                                                                                                                                      |
| block-size             | 4K                                                                                                                                                                                                                                                                                                                                                          |
| capacity               | 2G                                                                                                                                                                                                                                                                                                                                                          |
| description            | -                                                                                                                                                                                                                                                                                                                                                           |
| free-chunks            | ['0-524287']                                                                                                                                                                                                                                                                                                                                                |
| health-indications     | []                                                                                                                                                                                                                                                                                                                                                          |
| health-state           | ok                                                                                                                                                                                                                                                                                                                                                          |
| io-status              | alive                                                                                                                                                                                                                                                                                                                                                       |
| itls                   | 0x5000144230354911/0x5006016930600523/6,<br>0x5000144230354910/0x5006016930600523/6,<br>0x5000144230354910/0x5006016830600523/6,<br>0x5000144230354911/0x5006016830600523/6,<br>0x5000144220354910/0x5006016930600523/6,<br>0x5000144220354910/0x5006016830600523/6,<br>0x5000144220354911/0x5006016930600523/6,<br>0x5000144220354911/0x5006016830600523/6 |
| largest-free-chunk     | 2G                                                                                                                                                                                                                                                                                                                                                          |
| locality               | -                                                                                                                                                                                                                                                                                                                                                           |
| operational-status     | ok                                                                                                                                                                                                                                                                                                                                                          |
| storage-array-name     | EMC-CLARiion-APM00042201310                                                                                                                                                                                                                                                                                                                                 |
| storage-volumetype     | normal                                                                                                                                                                                                                                                                                                                                                      |
| system-id              | VPD83T3:6006016091c50e005057534d0c17e011                                                                                                                                                                                                                                                                                                                    |
| <b>thin-rebuild</b>    | <b>true</b>                                                                                                                                                                                                                                                                                                                                                 |
| total-free-space       | 2G                                                                                                                                                                                                                                                                                                                                                          |
| use                    | claimed                                                                                                                                                                                                                                                                                                                                                     |
| used-by                | []                                                                                                                                                                                                                                                                                                                                                          |
| vendor-specific-name   | DGC                                                                                                                                                                                                                                                                                                                                                         |

**Example** Claim multiple storage volumes whose names begin with VPD83T3:600601602:

```
Vplexcli:/clusters/cluster-1> storage-volume claim --storage-volumes VPD83T3:600601602*
```

- See also**
- ◆ [set on page 444](#)
  - ◆ [storage-volume claimingwizard on page 461](#)
  - ◆ [storage-volume unclaim on page 474](#)

## storage-volume claimingwizard

Finds unclaimed storage volumes, claims them, and names them appropriately.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **claimingwizard**.

**Syntax**

```
storage-volume claimingwizard
  [-c|--cluster] cluster
  [-f|--file] file,file...
  [-d|--dryRun]
  [-t|--set-tier] list
  --appc
  --thin-rebuild
```

**Arguments** **Required arguments**

None.

### Optional arguments

**[-c|--cluster] cluster** - Cluster on which to claim storage.

**[-f|--file] file,file...** - List of one or more files containing hints for storage-volume naming, separated by commas. Required for claiming volumes on storage arrays that do not include their array and serial number in response to SCSI inquiries.

**[-d|--dryRun]** - Do a dry-run only, do not claim and name the storage volumes.

**[-t|--set-tier] list** - Set a storage tier identifier per storage array in the storage-volume names. Type multiple `<arrayName>`, `<tier-character>` pairs separated by commas. Storage tier identifiers cannot contain underscores.

**--appc** - Make the specified storage volumes 'application consistent'. Prevents data already on the specified storage volume from being deleted or overwritten.



### CAUTION

**Once set, the application consistent attribute cannot be changed. This attribute can only be set when the storage- volumes/extents are in the claimed state.**

**--thin-rebuild** - Claims the specified storage volumes as "thin". Thin storage allocates blocks of data on demand versus allocating all the blocks up front. Thin provisioning eliminates almost all unused storage and improves utilization rates.

**Description** Storage volumes must be claimed, and optionally named before they can be used in a VPLEX cluster.

Storage tiers allow the administrator to manage arrays based on price, performance, capacity and other attributes. If a tier ID is assigned, the storage with a specified tier ID can be managed as a single unit. Storage volumes without a tier assignment are assigned a value of 'no tier'.

Use the **--set-tier** argument to add or change a storage tier identifier in the storage-volume names from a given storage array. For example:

```
Vplexcli:/clusters/cluster-1> storage-volume claimingwizard --set-tier="(Clar0400, L),
(Symm04A1, H)"
```

names all storage volumes from the CLARiiON array as `Clar0400L_<lun name>`, and all storage volumes from the Symmetrix® array as `Symm04A1H_<lun name>`

EMC Symmetrix, HDS 9970/9980 and USP V storage arrays include their array and serial number in response to SCSI inquiries. The claiming wizard can claim their storage volumes without additional information. Names are assigned automatically.

Other storage arrays require a hints file generated by the storage administrator using the array's command line. The hints file contains the device names and their World Wide Names.

Use the `--file` argument to specify a hints file to use for naming claimed storage volumes.

The following table lists examples to create hint files:

**Table 21 Create hints files for storage-volume naming**

| Storage array | Command to create hints file                                                                                                                                                                                                                                                                           |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EMC CLARiiON  | <code>navicli -h 192.168.47.27 getlun -uid -name &gt; Clar0400.txt</code>                                                                                                                                                                                                                              |
| EMC Symmetrix | <code>symdev -sid 781 list -wwn &gt; Symm0781.txt</code>                                                                                                                                                                                                                                               |
| EMC VPLEX     | <code>export view map -f EMC_PROD12.txt -v &lt;views&gt;</code>                                                                                                                                                                                                                                        |
| IBM DS4300    | <code>SMcli 192.168.97.121 -c "show logicalDrives;" &gt; DS4300_121.txt</code>                                                                                                                                                                                                                         |
| IBM Nextra    | <code>xcli -c nextra_lab -x vol_list &gt; Nextra_lab.txt</code>                                                                                                                                                                                                                                        |
| HP EVA        | <code>sssu "select manager &lt;hostname&gt; username=&lt;username&gt; password=&lt;password&gt;" "select system &lt;systemname&gt; "ls vdisk full" &gt; &lt;filename.txt&gt;</code>                                                                                                                    |
| Generic       | Text file of the following format:<br>> Generic storage-volumes<br>> VPD83T3:600a0b800011ea0a000073c5468cedbd MyName1<br>> 600a0b800011ea0a000073c5468cedbc MyName2<br>> vpd83t3:600A0b800011EA0a000073c5468cEdbd MyName3<br>For generic storage volumes, names may include letters, numbers, and '.'. |

**Example** In the following example, the `claimingwizard` command with no arguments claims storage volumes from an EMC Symmetrix array:

```
VPlexcli:/clusters/cluster-1> storage-volume claimingwizard
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> ll
```

| Name          | VPD83 ID                                 | Capacity | Use     | Vendor | IO Status | Type   | Thin  |
|---------------|------------------------------------------|----------|---------|--------|-----------|--------|-------|
| Rebuild       |                                          |          |         |        |           |        |       |
| -             |                                          |          |         |        |           |        |       |
| Symm2773_05F3 | VPD83T3:60000970000192602773533030354633 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F4 | VPD83T3:60000970000192602773533030354634 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F5 | VPD83T3:60000970000192602773533030354635 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F6 | VPD83T3:60000970000192602773533030354636 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F7 | VPD83T3:60000970000192602773533030354637 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F8 | VPD83T3:60000970000192602773533030354638 | 10G      | claimed | EMC    | alive     | normal | false |
| Symm2773_05F9 | VPD83T3:60000970000192602773533030354639 | 10G      | claimed | EMC    | alive     | normal | false |

Note that the Symmetrix storage volumes are named in the format:

`Symm<last 4 digits of array serial number>_<Symmetrix Device Number>`

**Example** In the following example:

- ◆ The `--cluster` argument specifies `cluster-1`

- ◆ The **--file** argument specifies a CLARiiON hints file containing device names and World Wide Names
- ◆ The **--thin-rebuild** argument claims the specified storage volumes as “thin” (data will be allocated on demand versus up front)

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> claimingwizard --cluster
cluster-1 --file /home/service/clar.txt --thin-rebuild
Found unclaimed storage-volume VPD83T3:6006016091c50e004f57534d0c17e011 vendor DGC : claiming
and naming clar_LUN82.
```

```
Found unclaimed storage-volume VPD83T3:6006016091c50e005157534d0c17e011 vendor DGC : claiming
and naming clar_LUN84.
```

```
Claimed 2 storage-volumes in storage array clar
```

```
Claimed 2 storage-volumes in total.
```

**Example** Find and claim storage volumes on any array in cluster-1 that does not require a hints file from the `/clusters/cluster/storage-elements/storage-volumes` context:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> claimingwizard

Found unclaimed storage-volume VPD83T1:HITACHI R45150040023 vendor HITACHI : claiming and
naming HDS20816_0023.

Found unclaimed storage-volume VPD83T1:HITACHI R45150040024 vendor HITACHI : claiming and
naming HDS20816_0024.
.
.
.
Fri, 20 May 2011 16:38:14 +0000 Progress : 6/101 storage_volumes processed (6%).
.
.
.
Fri, 20 May 2011 16:38:14 +0000 Progress : 96/101 storage_volumes processed (96%).
.
.
.
Claimed 37 storage-volumes in storage array Symm0487

Claimed 64 storage-volumes in storage array HDS20816

Claimed 101 storage-volumes in total.
```

- See also**
- ◆ [storage-volume claim on page 458](#)
  - ◆ [storage-volume unclaim on page 474](#)

## storage-volume find-array

Searches storage arrays for the specified storage-volume(s).

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **find-array**.

**Syntax** `storage-volume find-array  
[-d|--storage-volumes] pattern`

**Arguments** **Required arguments**  
`[-d|--storage-volumes] pattern` - Storage volume pattern for which to search.

The pattern conforms to 'glob'.

The following pattern symbols are supported: \*, ?, [seq], [!seq].

**Description** Searches all the storage arrays in all clusters for the specified storage-volume(s).  
The search is case-sensitive.

### Example

```
VPLEXcli:/> storage-volume find-array *de11
Searching for *de11
Storage-volume: VPD83T3:6006016021d0250026b925ff60b5de11 is in:
/clusters/cluster-1/storage-elements/storage-arrays/FNM00093200108-0x050060160bce03506
Storage-volume: VPD83T3:6006016021d0250027b925ff60b5de11 is in:
/clusters/cluster-1/storage-elements/storage-arrays/FNM00093200108-0x050060160bce03506
Storage-volume: VPD83T3:6006016021d0250028b925ff60b5de11 is in:
/clusters/cluster-1/storage-elements/storage-arrays/FNM00093200108-0x050060160bce03506
Storage-volume: VPD83T3:6006016021d0250029b925ff60b5de11 is in:
/clusters/cluster-1/storage-elements/storage-arrays/FNM00093200108-0x050060160bce03506
Storage-volume: VPD83T3:6006016021d025002ab925ff60b5de11 is in:
/clusters/cluster-1/storage-elements/storage-arrays/FNM00093200108-0x050060160bce03506
.
.
.
```

**See also** ♦ [storage-volume claimingwizard on page 461](#)

## storage-volume forget

Tells the cluster that a storage-volume or a set of storage volumes are physically removed.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **forget**.

**Syntax**

```
storage-volume forget
  [-d|--storage-volume] context path
  [-a|--all-at-cluster] context path
  [-c|--cluster] context path
  [-i|--logical-unit-id] logical unit id
```

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-d|--storage-volume] context path` - \* Storage volume to forget.

`[-a|--all-at-cluster] context path` - Forget all unclaimed, unreachable storage volumes at the specified cluster.

`[-c|--cluster] context path` - Forget the given internal system ID at a given cluster.

`[-i|--logical-unit-id] logical unitid` - Logical Unit ID (displayed in the storage-array/logical-units context) to forget.

\* - argument is positional.

**Description** Storage Volumes can be remembered (appear in the context tree) even if a cluster is not currently in contact with them.

Use the **storage-volume forget** command to tell the cluster that unclaimed and unreachable storage volumes are not coming back and it is safe to forget them.

Forgotten storage volumes are removed from the context tree.

Use both the **--logical-unit-id** and **--cluster** arguments to forget the Logical Unit ID in the specified storage-array/logical-units context.

Use the **--storage-volume** argument to forget only the specified unclaimed and unreachable storage volume.

Use the **--all-at-cluster** argument to forget all unclaimed and unreachable storage volumes at the specified cluster.

**Example** Forget a storage volume from the  
`/clusters/cluster/storage-elements/storage-volumes` context:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> forget --storage-volume  
VPD83T3:6006016021d0250027b925ff60b5de11
```

Forget a storage volume from the root context:

```
Vplexcli:> storage-volume forget --storage-volume  
clusters/cluster-1/storage-elements/storage-volumes/VPD83T3:6006016021d0250027b925ff60b5de11
```

**See also** ♦ [storage-volume unclaim on page 474](#)

## storage-volume list-banished

Displays banished storage-volumes on a director.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **list-banished**.

**Syntax** `storage-volume list-banished  
[-n|--director] context path`

**Arguments** **Required arguments**  
`[-n|--director] context path` - \*The director whose banished storage-volume(s) to display.

**Description** See [Banished storage volumes \(LUNs\) on page 472](#).

Displays the names of storage-volumes that are currently banished for a given director.

**Example** In the following example; director-1-1-A has one banished storage-volume:

```
Vplexcli:/> storage-volume list-banished --director director-1-1-A  
There is 1 banished storage-volume on director 'director-1-1-A':  
Symm0487_0C1B
```

**See also**

- ◆ [storage-volume auto-unbanish-interval on page 456](#)
- ◆ [storage-volume unbanish on page 472](#)

## storage-volume resurrect

Resurrect the specified storage-volume(s).

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **resurrect**.

**Syntax**

```
storage-volume resurrect
  [-d|--storage-volume] context path,context path...
  [-f|--force]
```

**Arguments** **Required arguments**

`[-d|--storage-volume] context path,context path...` - List of one or more storage volume with dead I/O status to resurrect.

**Optional arguments**

`[-f|--force]` - Force the storage-volume resurrect and bypass the test.

**Description** Resurrects the specified dead storage volumes and tests the resurrected device before setting its state to healthy.

A storage-volume is declared dead:

- ◆ After VPLEX retries a failed I/O to the backend arrays 20 times without success.
- ◆ If the storage-volume is reachable but errors prevent the I/O from succeeding.

A storage volume declared hardware dead cannot be unclaimed or removed (forgotten). Use this command to resurrect the storage volume. After the storage volume is resurrected, it can be unclaimed and removed.



### **IMPORTANT**

**Fix the root cause before resurrecting a storage volume because the volume can be successfully resurrected only to go back to dead on the next I/O.**

**Note:** This command will not work if the storage volume is marked unreachable.

This command has no ill effects if issued for a healthy storage volume.

LUNs exported from storage arrays can disappear or display I/O errors for various reasons, including:

- ◆ Marked read-only during copies initiated by the storage array
- ◆ Unrecoverable device errors
- ◆ Snapshot activation/deactivation on the storage array
- ◆ An operator shrinks the size of a storage-volume, causing the VPLEX to refuse to do I/O to the storage-volume.
- ◆ 100% allocated thin pools
- ◆ Persistent reservation on storage volume
- ◆ Dropped frames due to bad cable or SFP

Dead storage volumes are indicated by:

- The **cluster summary** command shows degraded health-state and one or more unhealthy storage volumes. For example:

```
Vplexcli:/clusters/cluster-2/> cluster status
Cluster cluster-2
operational-status:      ok
transitioning-indications:
transitioning-progress:
health-state:            degraded
health-indications:      1 unhealthy Devices or storage-volumes
```

- ◆ The **storage-volume summary** command shows the I/O status of the volume as dead. For example:

```
Vplexcli:/> storage-volume summary
SUMMARY (cluster-1)
StorageVolume Name      IO Status  Operational Status  Health State
-----
dead_volume              dead       error                critical-failure
```

Symptom:  
Storage-volume is dead

**Example** Resurrect two storage volumes:

```
Vplexcli:/> storage-volume resurrect --storage-volumes Symm1852_BAC,Symm1852_BA8
```

- See also**
- ◆ [cluster status on page 89](#)
  - ◆ [storage-volume forget on page 465](#)
  - ◆ [storage-volume summary on page 469](#)

## storage-volume summary

Displays a list of a cluster's storage volumes.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **summary**.

**Syntax** `storage-volume summary  
[-c|--clusters] cluster,cluster...`

### Optional arguments

`[-c|--clusters] cluster,cluster...` - Displays storage volumes for only the specified cluster(s).

**Description** Displays a two-part summary for each cluster's storage volumes:

- ◆ I/O status, operational status, and health state for each unhealthy storage volume.
- ◆ Summary of health-state, vendor, use, and total capacity for the cluster.

Use the `-clusters` argument to restrict output to only the specified cluster(s).

If no argument is used, and the command is executed at or below a `/clusters/cluster` context, output is for the specified *cluster* only.

Otherwise, output is for all clusters.

**Example** Display summary for only cluster-1 on a VPLEX with unhealthy volumes:

```
Vplexcli:/> storage-volume summary --clusters cluster-1
StorageVolume Name   IO Status   Operational Status   Health State
-----
Log1723_154          unreachable error                 critical-failure
Log1852_154          unreachable error                 critical-failure
Meta1723_150         unreachable error                 critical-failure
Meta1852_150         unreachable error                 critical-failure
Symm1378_0150        unreachable error                 critical-failure
Symm1378_0154        unreachable error                 critical-failure
.
.
Storage-Volume Summary (no tier)
-----
Health                out-of-date      0
                    storage-volumes 981
                    unhealthy       966

Vendor                DGC              15
                    None             966

Use                   claimed          824
                    meta-data        1
                    unclaimed        11
                    unusable         143
                    used             2

Capacity              total            16T
```

**Example** Display summary for both clusters in a VPLEX with no unhealthy storage volumes:

```
Vplexcli:/> storage-volume summary
SUMMARY (cluster-1)
```

```

Storage-Volume Summary (no tier)
-----
Health                out-of-date          0
                    storage-volumes    2318
                    unhealthy          0

Vendor                EMC                  2318

Use                   claimed              2172
                    meta-data           2
                    used                144

Capacity              total                198T

SUMMARY (cluster-2)
Storage-Volume Summary (no tier)
-----

Health                out-of-date          0
                    storage-volumes    2318
                    unhealthy          0

Vendor                EMC                  2318

Use                   claimed              2172
                    meta-data           2
                    used                144

Capacity              total                198T

```

**Table 22 storage-volume summary field descriptions (1 of 2)**

| Field                                                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Health summary (displayed only for unhealthy storage volumes)</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Name                                                                 | Name of storage volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| I/O Status                                                           | <b>alive</b> - I/O is proceeding normally on the storage volume.<br><b>dead</b> - VPLEX has marked the storage volume as dead; I/O cannot proceed on the storage volume. This can happen when a certain number of I/Os to the storage volume fails.<br><b>unreachable</b> - The storage volume is unreachable.                                                                                                                                                                                                                                                                |
| Operational Status                                                   | <b>ok</b> - The storage volume is functioning normally.<br><b>degraded</b> - The storage volume may be out-of-date compared to its mirror. (This state applies only to a storage volume that is part of a RAID 1 Metadata Volume.)<br><b>unknown</b> - VPLEX cannot determine the storage volume's Operational state, or the state is invalid.<br><b>error</b> - VPLEX has marked the storage volume as hardware-dead.<br><b>starting</b> - The storage volume is not yet ready.<br><b>lost communication</b> - The storage volume is unreachable.                            |
| Health State                                                         | <b>degraded</b> - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).<br><b>ok</b> - The extent is functioning normally.<br><b>non-recoverable-error</b> - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device), and/or the Health state cannot be determined.<br><b>unknown</b> - VPLEX cannot determine the extent's Operational state, or the state is invalid.<br><b>critical failure</b> - VPLEX has marked the storage volume as hardware-dead. |
| <b>Storage-Volume Summary</b>                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| out-of-date                                                          | Of the total number of storage volumes on the cluster, the number that are out-of-date compared to their mirror.                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 22 storage-volume summary field descriptions (2 of 2)**

| Field           | Description                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------|
| storage-volumes | Total number of storage volumes on the cluster.                                                       |
| unhealthy       | Of the total number of storage volumes on the cluster, the number with health state that is not "ok". |
| Vendor          | Of the total number of storage volumes on the cluster, the number from the specified vendor.          |
| claimed         | Of the total number of storage volumes on the cluster, the number that are claimed.                   |
| meta-data       | Of the total number of storage volumes on the cluster, the number in use as meta-volumes.             |
| unclaimed       | Of the total number of storage volumes on the cluster, the number that are unclaimed.                 |
| used            | Of the total number of storage volumes on the cluster, the number that are in use.                    |
| Capacity        | Total capacity of all storage on the cluster.                                                         |

- See also**
- ◆ [ds summary on page 234](#)
  - ◆ [export port summary on page 248](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [extent summary on page 270](#)
  - ◆ [local-device summary on page 283](#)
  - ◆ [storage-volume resurrect on page 467](#)
  - ◆ [virtual-volume summary on page 510](#)

## storage-volume unbanish

Unbanishes a storage-volume on one or more directors.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **unbanish**.

**Syntax** `storage-volume unbanish  
[-n|--directors] context path,context path...  
[-d|--storage-volume] context path`

**Arguments** **Required arguments**  
`[-n|--directors] context path,context path,...` - \* The directors on which to unbanish the given storage-volume.

**Optional arguments**  
`[-d|--storage-volume]` - The storage-volume to unbanish.

This argument is not required if the current context is a storage-volume or below. If the current context is a storage-volume or below, it operates on that storage volume.

\* - argument is positional.

**Description** VPLEX examines path state information for LUNs on arrays. If the path state information is inconsistent, VPLEX banishes the LUN, and makes it inaccessible.

Use this command to unbanish a banished LUN (storage volume).

### Banished storage volumes (LUNs)

LUNs (storage volumes) are banished when VPLEX detects an unexpected configuration of array controllers and/or paths to arrays. Under normal active/passive operation, one controller for any given LUN is active, the other is passive.

If the path to the active controller fails, the passive path transitions to active. The transition must wait for the failed active controller to drain its pending I/Os. This transient state may be seen during disk replacement, hot sparing, and disk failure.

If VPLEX detects a LUN in this state, VPLEX waits 20 seconds for the LUN to return to normal. If the LUN does not return to the expected state, VPLEX banishes the LUN.

**Example** In the following example:

- ◆ The **list-banished** command shows a volume is banished from director 1-1-A
- ◆ The **unbanish** command unbanishes the volume.
- ◆ The **list-banished** command shows the change:

```
Vplexcli:/> storage-volume list-banished --director director-1-1-A
There is 1 banished storage-volume on director 'director-1-1-A':
Symm0487_0C1B
```

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes/Symm0487_0C1B> storage-volume
unbanish --director director-1-1-A
director-1-1-A Unbanished.
```

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes/Symm0487_0C1B> storage-volume
list-banished --director director-1-1-A
There are no banished storage-volumes on director 'director-1-1-A'.
```

- 
- See also**
- ◆ [storage-volume auto-unbanish-interval on page 456](#)
  - ◆ [storage-volume list-banished on page 466](#)

## storage-volume unclaim

Unclaims the specified previously claimed storage volumes.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **unclaim**.

**Syntax** `storage-volume unclaim  
[-b|--batch-size] integer  
[-d|--storage-volumes] context path,context path...`

**Arguments** **Required arguments**  
`[-d|--storage-volumes] context path,context path...` - \* List of one or more storage volumes to unclaim.

**Optional arguments**  
`[-b|--batch-size] integer` - When using wildcards to unclaim multiple volumes with one command, the maximum number of storage volumes to unclaim at once.

\* - argument is positional.

**Description** Use the **storage-volume unclaim** command to return the specified storage volume(s) to the unclaimed state.

The target storage volume must not be in use.

### Unclaim a “thin” storage volume

When a storage-volume is unclaimed, the thin-rebuild attribute is set to false.

**Note:** The thin-rebuild attribute can only be modified for storage volumes that are either claimed or used. When the unclaimed storage-volume is claimed and its state is "claimed" or "used", use the **“set”** command to modify the thin-rebuild attribute.

**Example** In the following example:

- ◆ The **ll** command in storage-volumes context displays storage volumes, including their use state,
- ◆ The **storage-volume unclaim** command unclaims two claimed volumes:

```
VPlexcli:/clusters/cluster-2/storage-elements/storage-volumes> ll
```

| Name                        | VPD83 ID                                 | Capacity | Use     | Vendor | IO    | Type   | Thin Rebuild |
|-----------------------------|------------------------------------------|----------|---------|--------|-------|--------|--------------|
| Basic_c1_ramdisk_100GB_684_ | VPD83T3:60001440000000103017dfea88355431 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_685_ | VPD83T3:60001440000000103017dfea88355433 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_686_ | VPD83T3:60001440000000103017dfea88355435 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_687_ | VPD83T3:60001440000000103017dfea88355437 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_688_ | VPD83T3:60001440000000103017dfea88355439 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_689_ | VPD83T3:60001440000000103017dfea8835543b | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_68_  | VPD83T3:60001440000000103017dfea88354f61 | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_690_ | VPD83T3:60001440000000103017dfea8835543d | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_691_ | VPD83T3:60001440000000103017dfea8835543f | 100G     | claimed | EMC    | alive | normal | false        |
| Basic_c1_ramdisk_100GB_692_ | VPD83T3:60001440000000103017dfea88355441 | 100G     | claimed | EMC    | alive | normal | false        |

```
VPlexcli:/clusters/cluster-2/storage-elements/storage-volumes> unclaim -d  
Basic_c1_ramdisk_100GB_686_
```

- See also**
- ◆ [set on page 444](#)
  - ◆ [storage-volume claim on page 458](#)

---

## storage-volume used-by

Displays the components that use the specified storage volumes.

**Contexts** All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is **used-by**.

**Syntax** `storage-volume used-by  
[-d|--storage-volumes] context path,context path...`

**Arguments** **Required arguments**

`[-d|--storage-volumes] context path,context path...` - \* List of one or more storage volumes for which to find users.

**Description** To manually deconstruct an encapsulated storage volume, remove each layer starting from the top.

Use the **storage-volume used-by** command to see the layers from the bottom up.

**Example**

```
Vplexcli:/clusters/cluster-2/storage-elements/storage-volumes>  
used-by CX4_lun0  
/clusters/cluster-1/devices/base0:  
  extent_CX4_lun0_1  
  CX4_lun0  
  
/clusters/cluster-1/devices/base1:  
  extent_CX4_lun0_2  
  CX4_lun0  
  
/clusters/cluster-1/devices/base2:  
  extent_CX4_lun0_3  
  CX4_lun0  
  
/clusters/cluster-1/devices/base3:  
  extent_CX4_lun0_4  
  CX4_lun0  
  
/clusters/cluster-1/storage-elements/extents/extent_CX4_lun0_5:  
  CX4_lun0  
  
/clusters/cluster-1/storage-elements/extents/extent_CX4_lun0_6:  
  CX4_lun0
```

## subnet clear

Clears one or more attributes of an existing subnet configuration.

**Contexts** All contexts.

**Syntax**

```
subnet clear
  [-s|--subnet] subnet
  [-a|--attribute] attribute
```

**Arguments** **Required arguments**

**[-s|--subnet] *subnet*** - \* The subnet configuration to modify.

**Optional arguments**

**[-a|--attribute] *attribute*** - \* The name of the attribute to modify.

\* argument is positional

**Description** Clears (sets to null) one or more of the following attributes of an existing subnet configuration:

- ◆ cluster-address
- ◆ gateway
- ◆ prefix



### **IMPORTANT**

**Use the “[subnet modify](#)” command to configure new values for the attributes.**

**Example** In the following example:

- ◆ The **ll** command displays a subnet’s attributes.
- ◆ The **subnet clear** command clears (sets to null) the cluster-address attribute.
- ◆ The **subnet modify** command configures a new cluster-address.
- ◆ The **ll** command displays the new configuration:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> ll
Name                               Value
-----
cluster-address                    192.168.12.200
gateway                            192.168.12.1
mtu                                 9000
prefix                             192.168.12.0:255.255.255.0
proxy-external-address             -
remote-subnet-address              192.168.22.0:255.255.255.0
```

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> clear
cluster-address
```

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> ll
Name                               Value
-----
cluster-address                    -
gateway                            192.168.12.1
mtu                                 9000
prefix                             192.168.12.0:255.255.255.0
proxy-external-address             -
remote-subnet-address              192.168.22.0:255.255.255.0
```

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> subnet modify  
cluster-address 192.168.10.200
```

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> ll
```

| Name                   | Value                      |
|------------------------|----------------------------|
| cluster-address        | 192.168.10.200             |
| gateway                | 192.168.12.1               |
| mtu                    | 9000                       |
| prefix                 | 192.168.12.0:255.255.255.0 |
| proxy-external-address | -                          |
| remote-subnet-address  | 192.168.22.0:255.255.255.0 |

- See also**
- ◆ [subnet create on page 478](#)
  - ◆ [subnet destroy on page 480](#)
  - ◆ [subnet modify on page 481](#)

## subnet create

Creates a new subnet.

**Contexts** All contexts.

In /clusters/cluster-1/cluster-connectivity/subnets context, command is **create**.

**Syntax**

```
subnet create
  [-n|--name] subnet name
  [-c|--cluster] cluster
  [-a|--cluster-address] address
  [-g|--gateway] IP address
  [-m|--mtu] size
  [-p|--prefix] prefix
  [-e|--proxy-external-address] address
  [-r|--remote-subnet-address] remote subnet address
```

**Arguments** **Required arguments**

**[-n|--name] *subnet name*** - \* The name of the subnet to create.

**Optional arguments**

**[-c|--cluster] *cluster*** - Context path of the target owner cluster.

**[-a|--cluster-address] *address*** - The public address of the cluster to which this subnet belongs.

**[-g|--gateway] *IP address*** - The gateway address for this subnet.

**[-m|--mtu] *size*** - The maximum Transfer Unit size for this subnet.

**Range** - An integer between 96 and 9000.



### CAUTION

**The VPLEX CLI accepts MTU values lower than 96, but they are not supported. Entering a value less than 96 prevents the port-group from operating.**

**[-p|--prefix] *prefix*** - The prefix/subnet mask for this subnet. Specified as an IP address and subnet mask in integer dot notation, separated by a colon. For example, 192.168.20.0/255.255.255.0

**[-e|--proxy-external-address] *address*** - The externally published cluster address for this subnet. Can be one of:

- ◆ **w.x.y.z** where *w,x,y,z* are [0..255] - Configures the specified IP address. For example: 10.0.1.125.
- ◆ **empty** - Clears any configured IP address.

**[-r|--remote-subnet-address] *address*** - The [destination IP]:[netmask] subnet in the remote cluster that is reachable from the local subnet. Can be one of:

- ◆ **w.x.y.z** where *w,x,y,z* are [0..255] - Configures the specified IP address. For example: 172.16.2.0/255.255.255.0.
- ◆ **empty** - Clears any configured IP address

\* argument is positional.

**Description** Creates a new subnet configuration.

**Example** Create a subnet from the cluster-connectivity/subnet context:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets> create --name TestSubNet --cluster cluster-1 --gateway 192.168.10.1 --prefix 192.168.10.0/255.255.255.0
```

Subnet was created at '/clusters/cluster-1/cluster-connectivity/subnets/TestSubNet'.

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets> ll TestSubNet
```

```
/clusters/cluster-1/cluster-connectivity/subnets/TestSubNet:
```

| Name                   | Value                      |
|------------------------|----------------------------|
| -----                  | -----                      |
| cluster-address        | -                          |
| gateway                | 192.168.10.1               |
| mtu                    | 1500                       |
| prefix                 | 192.168.10.0/255.255.255.0 |
| proxy-external-address | -                          |
| remote-subnet-address  | -                          |

- See also**
- ◆ [subnet clear on page 476](#)
  - ◆ [subnet destroy on page 480](#)
  - ◆ [subnet modify on page 481](#)

## subnet destroy

Destroys a subnet configuration.

**Contexts** All contexts.

**Syntax**

```
subnet destroy
  [-s|--subnet] subnet
  --force
```

**Arguments** **Required arguments**

`[-s|--subnet] subnet` - \* Context path of the subnet configuration to destroy.

**Optional arguments**

`--force` - Forces the subnet configuration to be destroyed without asking for confirmation. Allows this command to be run from non-interactive scripts.

\* argument is positional

**Description** Destroys the specified subnet configuration.

This command fails if the specified subnet is currently assigned to any port or port-group.

**Example** Destroy a subnet from the /cluster-connectivity/subnets context:

```
VPlexcli:/clusters/cluster-2/cluster-connectivity/subnets> subnet destroy --subnet TestSubnet
```

**Example** Destroy a subnet from the root context:

```
VPlexcli:> subnet destroy /clusters/cluster-2/cluster-connectivity/subnets/TestSubnet
```

**See also**

- ◆ [subnet clear on page 476](#)
- ◆ [subnet create on page 478](#)
- ◆ [subnet modify on page 481](#)

## subnet modify

Modifies an existing subnet configuration.

**Contexts** All contexts.

**Syntax**

```
subnet modify
  [-s|--subnet] subnet
  [-a|--cluster-address] address
  [-g|--gateway] IP address
  [-p|--prefix] prefix
```

**Arguments** **Required arguments**

**[-s|--subnet] *subnet*** - \* Context path of the subnet configuration to modify.

**Optional arguments**

**[-a|--cluster-address] *address*** - The public address of the cluster to which this subnet belongs.

**[-g|--gateway] *IP address*** - The gateway address for this subnet.

**[-p|--prefix] *prefix*** - The prefix/subnet mask for this subnet. Specified as an IP address and subnet mask in integer dot notation, separated by a colon. For example, 192.168.20.0/255.255.255.0



### CAUTION

**If the prefix is changed, ensure that the cluster IP address, gateway address, and port IP addresses are all consistent with the subnet prefix.**

\* argument is positional.

**Description** Modifies one or more of the following attributes of an existing subnet configuration:

- ◆ cluster-address
- ◆ gateway
- ◆ prefix

**Example** Modify a subnet's prefix:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> subnet modify
--prefix 192.168.20.0/255.255.255.0
```

**Example** Modify a subnet's public IP address:

```
Vplexcli:/clusters/cluster-1/cluster-connectivity/subnets/cluster-1-SN01> subnet modify
--cluster-address 192.168.12.200
```

**See also**

- ◆ [subnet clear on page 476](#)
- ◆ [subnet destroy on page 480](#)
- ◆ [subnet create on page 478](#)

## syrcollect

Collects system configuration data for System Reporting (SYR).

**Contexts** All contexts.

**Syntax** `syrcollect  
[-d|--directory] directory`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-d|--directory] directory` - Non-default directory in which to store the output. Files saved in the non-default directory are not automatically sent to EMC.

Default: Files are stored in the Event\_Msg\_Folder in the directory specified in the EmaAdaptorConfig.properties file.

EmaAdaptorConfig.properties and the Event\_Msg\_Folder are located in /opt/emc/VPlex on the management server.

Files in the default directory are automatically sent to EMC.

**Description** Manually starts a collection of SYR data, and optionally sends the resulting zip file to EMC.

Run this command after every major configuration change or upgrade.

Data collected includes:

- ◆ VPLEX information
- ◆ RecoverPoint information (if RecoverPoint is configured)
- ◆ Cluster information
- ◆ Engine/chassis information
- ◆ RAID information
- ◆ Port information
- ◆ Back end storage information

The output of the command is a zipped xml file named:

`<VPLEXTLA>_Config_<TimeStamp>.zip`.

in the specified output directory.

Files in the default directory are automatically sent to EMC.

Use the **--directory** argument to specify a non-default directory. Output files sent to a non-default directory are not automatically sent to EMC.

**Example** Start a SYR data collection, and send the output to EMC:

```
Vplexcli:/> syrcollect
```

**Example** Start a SYR data collection, and send the output to the specified directory:

```
Vplexcli:/> syrcollect -d /var/log/VPlex/cli
```

**See also**

- ◆ [scheduleSYR add on page 413](#)
- ◆ [scheduleSYR list on page 414](#)

- ◆ `scheduleSYR remove` on page 415

## tree

Displays the context tree.

**Contexts** All contexts.

**Syntax**

```
tree
  [-e|--expand]
  [-c|--context] subcontext root
  [-s|--select] glob pattern
```

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-e|--expand]** - Expand the subcontexts.

**[-c|--context] *subcontext root*** - The subcontext to use as the root for the tree.

**[-s|--select] *glob pattern*** - Glob pattern for selecting the contexts in the tree.

**Description** Displays the sub-context tree.

Use the **tree** command with no arguments to display the sub context tree from the current context.

Use the **--context *subcontext root*** to display the sub context tree from the specified subcontext.

Use the **--expand** argument to expand the sub-contexts if applicable.

Use the **--select *glob pattern*** argument to display contexts in the specified sub-tree that match the glob pattern. The glob pattern may match more contexts that are outside the given sub-tree.

**Example** Display contexts below the current context:

```
Vplexcli:/management-server> tree
/management-server:
  ports
    eth0
    eth1
    eth2
    eth3
```

Display contexts below the specified context:

```
Vplexcli:/> tree --context /clusters/cluster-1/devices/dev_sym1723_1FC
/clusters/cluster-1/devices/dev_sym1723_1FC:
  components
    extent_Symm1723_1FC_1
      components
        Symm1723_1FC
          components
```

**See also** ♦ [drill-down on page 217](#)

## unalias

Removes a command alias.

**Contexts** All contexts.

**Syntax**

```
unalias
  [-n|--name] name
  [-a|--all]
```

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-n|--name] name** - The name of the alias to remove.

**[-a|--all]** - Remove all defined aliases.

**Example** In the following example:

- ◆ **alias** displays a list of all aliases on the VPLEX
- ◆ **unalias** deletes the specified alias
- ◆ **alias** confirms the deletion:

```
Vplexcli:/> alias
Name          Description
-----
?             Substitutes the 'help' command.
GoToDir_2_2A  Substitutes the 'cd
               /engines/engine-2-2/directors/Cluster_2_Dir_2A' command.
ll            Substitutes the 'ls -al' command.
quit         Substitutes the 'exit' command.
```

```
Vplexcli:/> unalias GoToDir_2_2A
```

```
Vplexcli:/> alias
Name  Description
-----
?     Substitutes the 'help' command.
ll    Substitutes the 'ls -al' command.
quit  Substitutes the 'exit' command.
```

**See also** ◆ [alias on page 27](#)

## user add

Adds a username to the VPLEX management server.

**Contexts** All contexts.

**Syntax**  
`user add [-u|--username] username`

**Arguments** **Required arguments**  
`[-u|--username] username` - Username to add.

**Description** Administrator privileges are required to execute the “`user add`” command.  
VPLEX has two pre-configured CLI users that can not be removed: **admin** and **service**.

---

**Note:** In VPLEX Metro and Geo configuration, VPLEX CLI accounts created on one management server are not propagated to the second management server. The **user list** command displays only those accounts configured on the local management server, not both server.

---

**Note:** Administrative privileges are required to add, delete, and reset user accounts. The password for the admin account must be reset the first time the admin account is accessed. After the admin password has been reset, the admin user can manage (add, delete, reset) user accounts.

To change the password for the admin account, ssh to the management server as user “admin”. Enter the default password `teS6nAX2`. A prompt to change the admin account password appears. Enter a new password.

---

**Example** Login to the CLI as an Administrator user.

At the CLI prompt, type the **user add *username*** command:

```
VPlxcli: /> user add TestUser
```

A prompt for the Administrator password appears:

```
admin password:
```

Type the password for the Administrator username.

A prompt for the new password for the username being added appears:

```
New password:
```

Type the password for the new username.

A prompt to confirm the password appears:

```
Confirm password:
```

Re-type the password.

Type the **user list** command to verify the new username is added:

```
VPlxcli: /> user list  
Username  
-----  
admin
```

```
service
TestUser
```

```
Vplexcli: />
```

- See also**
- ◆ “user list” on page 488
  - ◆ “user passwd” on page 489
  - ◆ “user remove” on page 490
  - ◆ “user reset” on page 491

## user list

Displays usernames configured on the local VPLEX management server.

**Contexts** All contexts.

**Syntax** `user list`

**Description** Displays the configured usernames.

---

**Note:** In VPLEX Metro and Geo configuration, VPLEX CLI accounts created on one management server are not propagated to the second management server. The **user list** command displays only those accounts configured on the local management server, not both servers.

---

**Example** Display the user accounts configured on the local management server:

```
Vplexcli: /> user list
Username
-----
admin
service
TestUser
```

**See also**

- ◆ [“user add” on page 486](#)
- ◆ [“user passwd” on page 489](#)
- ◆ [“user remove” on page 490](#)
- ◆ [“user reset” on page 491](#)

---

## user passwd

Allows a user to change the password for their own username.

**Contexts** All contexts.

**Syntax** `user passwd  
[-u|--username] username`

**Arguments** **Required arguments**  
`[-u|--username] username` - \*Username for which to change the password.  
\* - argument is positional.

**Description** Executable by all users to change the password only for their own username.

**Example** `Vplexcli: /> user passwd TestUser`

A prompt for the current password appears:

old password:

Type the current password for the username.

A prompt for the new password appears:

New password:

Type the new password. Passwords must be at least 8 characters long, and must not be dictionary words.

A prompt to confirm the new password appears:

Confirm password:

Retype the new password.

**See also**

- ◆ [“user add” on page 486](#)
- ◆ [“user list” on page 488](#)
- ◆ [“user remove” on page 490](#)
- ◆ [“user reset” on page 491](#)

---

## user remove

Removes a username from the VPLEX management server.

**Contexts** All contexts.

**Syntax**  
`user remove  
[-u|--username] username`

**Arguments** **Required arguments**  
`[-u|--username] username` - Username to remove.

**Description** Administrator privileges are required to execute the “[user remove](#)” command.

---

**Note:** Administrative privileges are required to add, delete, and reset user accounts. The password for the admin account must be reset the first time the admin account is accessed. After the admin password has been reset, the admin user can manage (add, delete, reset) user accounts.

To change the password for the admin account, ssh to the management server as user “admin”. Enter the default password teS6nAX2. A prompt to change the admin account password appears. Enter a new password.

---

**Example** Login as an Administrator user.

Type the `user remove username` command:

```
Vplexcli: /> user remove TestUser
```

A prompt for the Administrator password appears:

```
admin password:
```

Type the password for the Administrator username.

The specified user username is removed.

Type the “[user list](#)” command to verify the removal.

**See also**

- ◆ [“user add” on page 486](#)
- ◆ [“user list” on page 488](#)
- ◆ [“user passwd” on page 489](#)
- ◆ [“user reset” on page 491](#)

---

## user reset

Allows an Administrator user to reset the password for any username.

**Contexts** All contexts.

**Syntax**  
`user reset  
[-u|--username] username`

**Arguments** **Required arguments**  
`[-u|--username] username` - The username whose password is to be reset.

**Description** Resets the password for any username.  
Administrator privileges are required.

---

**Note:** Administrative privileges are required to add, delete, and reset user accounts. The password for the admin account must be reset the first time the admin account is accessed. After the admin password has been reset, the admin user can manage (add, delete, reset) user accounts.

To change the password for the admin account, ssh to the management server as user "admin". Enter the default password teS6nAX2. A prompt to change the admin account password appears. Enter a new password.

All users can change the password for their own account using the "user passwd" command.

**Example** Login as an Administrator user.

Type the **user reset --username *username*** command:

```
Vplexcli: /> user reset --username TestUser
```

A prompt for the Administrator password appears:

```
admin password:
```

Type the password for the Administrator username.

A prompt for new password for the username being reset appears:

```
New password:
```

Type a new password for the username.

A prompt to confirm the new password appears:

```
Confirm password:
```

Re-type the new password.

**See also**

- ◆ "user add" on page 486
- ◆ "user list" on page 488
- ◆ "user remove" on page 490

## validate-system-configuration

Performs a basic system configuration check.

**Contexts** All contexts.

**Syntax** validate-system-configuration

**Description** This command performs the following checks:

- ◆ Validates cache mirroring.
- ◆ Validates the logging volume.
- ◆ Validates the meta-volume.
- ◆ Validates back-end connectivity.

**Example** Validate system configuration:

```
Vplexcli:/> validate-system-configuration
Validate cache replication
Checking cluster cluster-1 ...
rmg component not found skipping the validation of cache replication.
ok
```

```
Validate logging volume
No errors found
ok
```

```
Validate back-end connectivity
Cluster cluster-2
  0 storage-volumes which are dead or unreachable.
  0 storage-volumes which do not meet the high availability
  requirement for storage volume paths*.
  0 storage-volumes which are not visible from all directors.
```

\*To meet the high availability requirement for storage volume paths each storage volume must be accessible from each of the directors through 2 or more VPLEX backend ports, and 2 or more Array target ports, and there should be 2 or more ITLs.

```
Cluster cluster-1
  10 storage-volumes which are dead or unreachable.
  0 storage-volumes which do not meet the high availability
  requirement for storage volume paths*.
  0 storage-volumes which are not visible from all directors.
```

\*To meet the high availability requirement for storage volume paths each storage volume must be accessible from each of the directors through 2 or more VPLEX backend ports, and 2 or more Array target ports, and there should be 2 or more ITLs.

Errors were encountered in the back-end connectivity. Please run 'connectivity validate-be -s' for details.

```
Validate meta-volume
Checking cluster cluster-1 ...
Checking cluster cluster-2 ...
ok
```

- See also**
- ◆ [cluster status on page 89](#)
  - ◆ [connectivity validate-be on page 142](#)
  - ◆ [health-check on page 273](#)

---

## vault go

Initiates a manual vault on every director in a given cluster under emergency conditions.

**Contexts** All contexts.

**Syntax**

```
vault go
  [-c|--cluster] cluster
  [--force]
```

**Arguments** [-c|--cluster] *cluster* - Specify the cluster on which to start cache vaulting.  
[--force] - Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.

**Description** Use this command to initiate a manual dump from every director in a given cluster to persistent local storage under emergency conditions.

Use this command to manually start cache vaulting if an emergency shutdown is required and the storage administrator cannot wait for automatic vaulting to begin.

**Example** Start a manual vault on cluster-1:

```
Vplexcli: /> vault go --cluster cluster-1
```

**See also**

- ◆ [“vault overrideUnvaultQuorum” on page 494](#)
- ◆ [“vault status” on page 496](#)
- ◆ *EMC VPLEX Administrator's Guide*

## vault overrideUnvaultQuorum

Allows the cluster to proceed with the recovery of the vault(s) without all the required director(s).

**Contexts** All contexts.

**Syntax**

```
vault overrideUnvaultQuorum
  [-c|--cluster] cluster
  --evaluate-override-before-execution
  --force
```

**Arguments** [-c|--cluster] *cluster* - Overrides unvault quorum for the specified cluster.

**--evaluate-override-before-execution** - Evaluates the possible outcome of running this command but does not do anything.

**--force** - Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.

### Description



#### **WARNING**

***This command could result in data loss.***

Use this command to tell the cluster not to wait for all the required director(s) to rejoin the cluster before proceeding with vault recovery.

Use this command with the **--evaluate-override-before-execution** argument to evaluate the cluster's vault status and make a decision whether to accept a possible data loss and continue to bring the cluster up. The evaluation provides information as to whether the cluster has sufficient vaults to proceed with the vault recovery that will not lead to data loss.

**Note:** One valid vault can be missing without experiencing data loss.

### Example Evaluate vault quorum

In the following example, the **--evaluate-override-before-execution** argument evaluates the cluster's unvault quorum state in the following circumstances:

- ◆ Three directors in a dual engine configuration booted and joined the cluster.
- ◆ None of these directors have a valid vault.
- ◆ The cluster is waiting for the remaining director to join the cluster before unvault recovery quorum is established.

```
Vplexcli:/> vault overrideUnvaultQuorum --evaluate-override-before-execution -c cluster-1
Cluster's unvault recovery quorum status:
```

```
Only 3 out of 4 configured directors on this cluster are running, and none has reported a
valid vault.
```

```
All configured directors must be present to verify if any director has successfully vaulted
dirty data the last time the cluster was servicing I/O.
```

```
Missing directors in the cluster: director-1-1-A
```

### Example Override unvault quorum wait state (no valid vaults)

In the following example, the command overrides the unvault quorum wait state in the following circumstances:

- ◆ None of the operational directors have valid vaults

- ◆ One director is not operational

```
VPLEXcli:/> vault overrideUnvaultQuorum -c cluster-1
Warning: Execution of this command can result in possible data loss based on the current vault
status of the cluster.
Cluster's unvault recovery quorum status:
Only 3 out of 4 configured directors on this cluster are running, and none has reported a valid
vault.
All configured directors must be present to verify if any director has successfully vaulted
dirty data the last time the cluster was servicing I/O.
    Missing directors in the cluster: director-1-1-A

Do you wish to override unvault quorum? (Yes/No) No

Execution of the override unvault quorum has been canceled by user!
```

**Example Override unvault quorum wait state (3 of 4 valid vaults)**

In the following example, the command overrides the unvault quorum wait state in the following circumstances:

- ◆ Three of four configured directors are operational and have valid vaults
- ◆ One director is not operational

```
VPLEXcli:/> vault overrideUnvaultQuorum -c cluster-1
Warning: Execution of this command can result in possible data loss based on the current vault
status of the cluster.
Cluster's unvault recovery quorum status:
3 out of 4 directors that were servicing I/O the last time the cluster had vaulted are present,
which is sufficient to proceed with vault recovery.

Do you wish to override unvault quorum? (Yes/No) Yes
Execution of the override unvault quorum has been issued!
```

**Example Override unvault quorum wait state (2 of 4 valid vaults)**

In the following example, the command evaluates the cluster vault status and overrides the unvault quorum when:

- ◆ Three of four configured directors are operational
- ◆ Two operational directors have valid vaults

```
VPLEXcli:/> vault overrideUnvaultQuorum -c cluster-1

Warning: Execution of this command can result in possible data loss based on the current vault
status of the cluster.

Cluster's unvault quorum status:
Only 2 out of 4 directors that were servicing I/O the last time the cluster had vaulted are
present, which is insufficient to proceed with vault recovery.
    Number of directors reporting a valid vault in the cluster: 2
    Number of directors reporting an invalid vault in the cluster: 1
    Missing directors in the cluster: director-1-2-B

Do you wish to override unvault quorum? (Yes/No) No

Execution of the override unvault quorum has been canceled by the user!!
```

- See also**
- ◆ [“vault status” on page 496](#)
  - ◆ [“vault go” on page 493](#)
  - ◆ *EMC VPLEX Administrator’s Guide*
  - ◆ VPLEX Procedure Generator

## vault status

Displays the current cache vault/unvault status of the cluster.

**Contexts** All contexts.

**Syntax**  
`vault status`  
`[-c|--cluster] cluster`

**Arguments** **Optional arguments**  
`[-c|--cluster] cluster` -Displays vault status for the specified cluster.  
`--verbose` - Displays additional description and data.

**Description** Cache vaulting safeguards dirty data during power outages. Cache vaulting dumps all dirty data to persistent local storage. Vaulted data is recovered (unvaulted) when power is restored.

This command always displays the cluster's vault state and the vault state of each of the cluster's directors.

When run after a vault has begun and the vault state is Vault Writing or Vault Written, the following information is displayed:

- ◆ Total number of bytes to be vaulted in the cluster
- ◆ Estimated time to completion for the vault

When run after the directors have booted and unvaulting has begun and the states are Unvaulting or Unvault Complete, the following information is displayed:

- ◆ Total number of bytes to be unvaulted in the cluster
- ◆ Estimated time to completion for the unvault
- ◆ Percent of bytes remaining to be unvaulted
- ◆ Number of bytes remaining to be unvaulted

If you enter the `--verbose` argument, the command displays the following information:

- ◆ Average vault or unvault rate.

If this command is run after the directors have booted, unvaulted, and are waiting to acquire an unvault quorum:

- ◆ The state is Unvault Quorum Waiting.
- ◆ The output displays a list of directors that are preventing the cluster from gaining unvault quorum.

If the `--verbose` argument is used, the following additional information is displayed:

- ◆ Vaulted data is valid or invalid.
- ◆ Required number of directors to proceed with the recovery of vault.

- ◆ Number of directors missing and preventing the cluster from proceeding with the recovery of vault.

**Table 23 Vault state field descriptions**

| Vault State            | Description                                                                                   |
|------------------------|-----------------------------------------------------------------------------------------------|
| Vault Inactive         | Vault/Unvault is not in progress.                                                             |
| Power Loss Detected    | Power loss has been detected. Waiting for power to be restored.                               |
| Power Loss Confirmed   | Power did not get restored.                                                                   |
| Power Restored         | Power has been restored.                                                                      |
| Vault Requested        | Initial vault request has been received.                                                      |
| Stop I/O               | The director is stopping all I/O.                                                             |
| Vault Writing          | Dirty data is being written to local persistent storage.                                      |
| Vault Written          | Dirty data has been written to local persistent storage.                                      |
| Unvaulting             | Vaulted dirty data is being read from local persistent storage.                               |
| Unvault Complete       | All vaulted dirty data has been read from local persistent storage.                           |
| Unvault Quorum Waiting | Waiting on all director(s) required before proceeding with the recovery of the vault.         |
| Vault Recovering       | Recovering all the vaulted dirty data to VPLEX global cache from the local director's memory. |

**Example** Display the summarized status for a cluster that is not currently vaulting or unvaulting:

```
Vplexcli: /> vault status --cluster cluster-1

=====
Cluster level vault status summary
=====
Cluster:/clusters/cluster-1
Cluster is not vaulting/unvaulting.
=====
Director level summary
=====
/engines/engine-1-1/directors/director-1-1-B:
state: Vault Inactive
/engines/engine-1-1/directors/director-1-1-A:
state: Vault Inactive
```

**Example** Display detailed status when a cluster is vaulting:

```
Vplexcli: /> vault status -c cluster-2 --verbose

=====
Cluster level vault status summary
=====
Cluster:/clusters/cluster-2
Cluster is vaulting
Total number of bytes remaining to vault in the cluster: 1.104809984 GB
Estimated time remaining for cluster's vault completion: 10 seconds
=====
Director level vault status summary
=====
/engines/engine-2-1/directors/director-2-1-B:
state: Vaulting - Writing vault data to vault disk
Total number of bytes to vault: 566.403072 MB
Total number of bytes vaulted: 8.220672 MB
Total number of bytes remaining to vault: 558.182400 MB
```

```

Percent vaulted: 1%
Average vault rate: 54.050816 MB/second
Estimated time remaining to vault complete: 10 seconds
/engines/engine-2-1/directors/director-2-1-A:
state: Vaulting - Writing vault data to vault disk
Total number of bytes to vault: 554.848256 MB
Total number of bytes vaulted: 8.220672 MB
Total number of bytes remaining to vault: 546.627584 MB
Percent vaulted: 1%
Average vault rate: 51.875840 MB/second
Estimated time remaining to vault complete: 10 seconds

```

**Example** Display detailed information when a cluster is un-vaulting:

```

VPlexcli:/> vault status -c cluster-2 --verbose

```

```

=====
Cluster level unvault status summary
=====
Cluster:/clusters/cluster-2
Cluster is unvaulting.
Total number of bytes remaining to unvault in the cluster: 583.499776 MB
Estimated time remaining for cluster's unvault completion: 24 seconds
=====
Director level unvault status summary
=====
/engines/engine-2-1/directors/director-2-1-B:
state: Unvaulting - Reading vaulted data from vault in to the local
director's memory
Total number of bytes to unvault: 566.403072 MB
Total number of bytes unvaulted: 289.505280 MB
Total number of bytes remaining to unvault: 276.897792 MB
Percent unvaulted: 51%
Average unvault rate: 14.471168 MB/second
Estimated time remaining to unvault complete: 19 seconds
/engines/engine-2-1/directors/director-2-1-A:
state: Unvaulting - Reading vaulted data from vault in to the local
director's memory
Total number of bytes to unvault: 554.848256 MB
Total number of bytes unvaulted: 248.246272 MB
Total number of bytes remaining to unvault: 306.601984 MB
Percent unvaulted: 44%
Average unvault rate: 12.410880 MB/second
Estimated time remaining to unvault complete: 24 seconds

```

**Example** Display when a cluster is waiting to acquire an unvault quorum:

```

VPlexcli:/> vault status -c cluster-1 --verbose

```

```

=====
Cluster level summary
=====
Cluster:/clusters/cluster-1
Cluster is waiting on all director(s) required before proceeding with the recovery of vault
Number of directors required to be present before cluster can proceed with the recovery
of the vault : 4
Number of directors present with valid vaults: 0
Number of directors present with invalid vaults: 3
Number of directors missing and possibly preventing the cluster to proceed with the
recovery of the vault : 1
Missing directors: director-1-1-A
=====
Director level summary
=====
/engines/engine-1-1/directors/director-1-1-B:

```

```
state: Waiting for unvault recovery quorum - Waiting on all director(s) required before
proceeding with the recovery of vault
Vault does not contain any data
Required number of directors to proceed with the recovery of vault: 4
Number of directors preventing the cluster to proceed with the recovery of vault: 1
Missing director list: director-1-1-A
/engines/engine-1-2/directors/director-1-2-B:
state: Waiting for unvault recovery quorum - Waiting on all director(s) required before
proceeding with the recovery of vault
Vault does not contain any data
Required number of directors to proceed with the recovery of vault: 4
Number of directors preventing the cluster to proceed with the recovery of vault: 1
Missing director list: director-1-1-A
/engines/engine-1-2/directors/director-1-2-A:
state: Waiting for unvault recovery quorum - Waiting on all director(s) required before
proceeding with the recovery of vault
Vault does not contain any data
Required number of directors to proceed with the recovery of vault: 4
Number of directors preventing the cluster to proceed with the recovery of vault: 1
Missing director list: director-1-1-A
/engines/engine-1-1/directors/director-1-1-A:
director could not be reached
```

- See also**
- ◆ [“vault go” on page 493](#)
  - ◆ [“vault overrideUnvaultQuorum” on page 494](#)
  - ◆ *EMC VPLEX Administrator's Guide*

## verify fibre-channel-switches

Verifies that the Fibre Channel switch on each cluster's internal management network has been configured correctly.

**Contexts** /clusters/cluster-n

**Syntax** verify fibre-channel-switches

Passwords for the service accounts on the switches are required to run this command.

**Example** Verify the internal management network:

```
VPLEXcli:/clusters/cluster-2> verify fiber-channel-switches
```

```
Please enter the service account password for the Management Server:
```

```
Re-enter the password:
```

```
Please enter the service account password for the fibre channel switch at IP 128.221.252.66:
```

```
Re-enter the password:
```

```
Executing on the first switch.  
This will take a few moments...
```

```
Please enter the service account password for the fibre channel switch at IP 128.221.253.66:
```

```
Re-enter the password:
```

```
Executing on the second switch.  
This will take a few moments...
```

```
Both fiber channel switch configurations were successfully verified.
```

## version

Display version information for connected directors.

**Contexts** All contexts.

**Syntax**  
version  
[-a|--all]  
[-n|directors] context path, context path...

**Arguments** **Required arguments**  
None.

**Optional arguments**

**[-a|--all]** - Displays version information for all connected directors.

**[-n|--directors] context path, context path...** - \* Display version information for only the specified directors.

\* - argument is positional.

**Description** Use the **--verbose** option to display version information for individual software components on each director.

**Example** Display management server/SMS version information:

| What                 | Version                  | Info |
|----------------------|--------------------------|------|
| Product Version      | 4.2.0.00.00.08           | -    |
| SMSv2                | D4.0.0.0-20101220.0      | -    |
| Mgmt Server Base     | D4_MSB_7                 | -    |
| Mgmt Server Software | D4.0.0.0.20101220000000e | -    |

**Example** Display management server/SMS version and version for the specified director:

```
Vplexcli:/> version Cluster_2_Dir_1B
```

| What                                           | Version     | Info |
|------------------------------------------------|-------------|------|
| SMSv2                                          | 0.16.15.0.0 | -    |
| Mgmt Server Base                               | D4_MSB_7    | -    |
| Mgmt Server Software                           | D4.70.0.9   | -    |
| /engines/engine-2-1/directors/Cluster_2_Dir_1B | 1.2.43.9.0  | -    |

**Example** Display version information for management server, SMS, and all directors:

```
Vplexcli:/> version -a
```

| What                                           | Version        | Info |
|------------------------------------------------|----------------|------|
| Product Version                                | 4.1.0.00.00.12 | -    |
| SMSv2                                          | 0.16.15.0.0    | -    |
| Mgmt Server Base                               | D4_MSB_7       | -    |
| Mgmt Server Software                           | D4.70.0.9      | -    |
| /engines/engine-2-1/directors/Cluster_2_Dir_1B | 1.2.43.9.0     | -    |
| /engines/engine-2-1/directors/Cluster_2_Dir_1A | 1.2.43.9.0     | -    |
| .                                              |                |      |
| .                                              |                |      |
| .                                              |                |      |

**Example** Display version information for individual software components on each director. See [Table 24](#) for a description of the components.

```
Vplexcli:/> version -a --verbose
```

```
Product Version: 5.1.0.00.00.08
```

What: SMSv2  
Version: D10.0.0.189.0  
Build time: March 29, 2012 at 08:23:02PM MDT  
Build machine: dudleyed06  
Build OS: Linux version 2.6.27-7-generic on i386  
Build compiler: 1.6.0  
Build source: /spgear/spgear\_misc/htdocs/harness/release/235/work/ui/src

What: Mgmt Server Base  
Version: D4\_MSB\_7

What: Mgmt Server Software  
Version: D10.0.0.212

What: Cluster Witness Server Software  
Version: D10.0.0.199

For director /engines/engine-2-2/directors/director-2-2-B:

What: O/S  
Version: D10.0.0.70 (SLES11)

What: Director Software  
Version: 2.1.191.0.0

What: ECOM  
Version: 4.1.13.0.0-0

What: VPLEX Splitter  
Version: 3.5.n.109-1

What: ZECL  
Version: 2.1.195.0.0-0

What: ZPEM  
Version: 2.1.195.0.0-0

What: NSFW  
Version: 40.1.160.0-0

What: FW Bundle Rev  
Version: 12.10 (Wed Apr 27 09:21:15 2011)

What: POST Rev  
Version: 41.00 (Mon Apr 25 13:39:19 2011)

What: BIOS Rev  
Version: 7.70 (05-11-2011)

What: SSD Model: P30056-MTFDBAA056SAL 118032803  
Version: 0005

.  
.  
.

**Table 24 Software components**

| <b>Component Abbreviation</b>   | <b>Description</b>                                                                                                                                                                                                                                                                                 |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Product Version                 | VPLEX version information.                                                                                                                                                                                                                                                                         |
| Mgmt Server Base                | Novel Linux distribution.                                                                                                                                                                                                                                                                          |
| Mgmt Server Software            | Version of the software on the management server.                                                                                                                                                                                                                                                  |
| Director Operating System       | Novell Linux distribution.                                                                                                                                                                                                                                                                         |
| Cluster Witness Server Software | Version of the Cluster Witness Sever VM.                                                                                                                                                                                                                                                           |
| OS                              | Operating system running on the director.                                                                                                                                                                                                                                                          |
| Director Software               | Version of the software on the specified director.                                                                                                                                                                                                                                                 |
| ECOM                            | The EMC Common Object Manager is a hub of communications and common services for applications based on EMC's Common Management Platform.                                                                                                                                                           |
| VPLEX Splitter                  | RecoverPoint splitter version.                                                                                                                                                                                                                                                                     |
| ZECL                            | A kernel module in the director that interfaces with the ZPEM process to provide, among other things, access to the I2C bus.                                                                                                                                                                       |
| ZPEM                            | Manages the overall health of the hardware. It includes monitoring of the various Field Replaceable Units (FRUs), Power and Temperature values and monitoring of external entities like the Standby Power Supply (SPS), COM FC Switch and the UPS used to provide backup power to the FC switches. |
| NSFW                            | VPLEX GeoSynchrony™ software. VPLEX operating system running in the cluster(s).                                                                                                                                                                                                                    |
| FW Bundle Rev                   | Firmware revision.                                                                                                                                                                                                                                                                                 |
| POST Rev                        | Power On Self Test revision.                                                                                                                                                                                                                                                                       |
| BIOS Rev                        | Boot firmware revision.                                                                                                                                                                                                                                                                            |
| SSD Model                       | Solid state disk drive model information.                                                                                                                                                                                                                                                          |

## virtual-volume create

Creates a virtual volume on a host device.

**Contexts** All contexts.

**Syntax**  
`virtual-volume create`  
`[-r|--device] context path`  
`[-t|--set-tier] tier]`

**Arguments** **Required arguments**  
`[-r|--device] context path` - \* Device on which to host the virtual volume.

**Optional arguments**  
`[-t|--set-tier]` - Set the storage-tier for the new virtual volume.

\* - argument is positional.

**Description** A virtual volume is created on a device or a distributed device, and is presented to a host through a storage view. Virtual volumes are created on top-level devices only, and always use the full capacity of the device or distributed device.

The underlying storage of a virtual volume may be distributed over multiple storage volumes, but appears as a single contiguous volume.

The specified device must not already have a virtual volume and must not have a parent device.

Use the `--set-tier` argument to set the storage tier for the new virtual volume.

### About storage tier IDs

Storage tiers are user-defined classes of storage, usually based on price, performance, capacity, and other factors. Storage tiers allow administrators to group like arrays so they can be managed as a unit.

The storage-tier identifier is displayed to the host as part of the virtual volumes's product ID.

Use the storage-tier identifier to logically group storage.

For example, assign Symmetrix arrays as tier 1 storage, and CLARiiON as tier 2 storage.

Use the `ll` command in a specific virtual volume's context to display the current storage-tier.

Use the `set` command to modify a virtual volume's storage-tier.

**Example** In the following example:

- ◆ The `virtual-volume create` command creates a new virtual volume,
- ◆ The `cd` command navigates to the new virtual volume's context,
- ◆ The `ll` command displays the new virtual volume:

```
Vplexcli:/> virtual-volume create --device
/distributed-storage/distributed-devices/TestDevice --set-tier 1
Vplexcli:/> cd /clusters/cluster-1/virtual-volumes/TestDevice_vol/
```

```
Vplexcli:/clusters/cluster-1/virtual-volumes/TestDevice_vol> ll
Name                               Value
-----
block-count                         264000
block-size                           4K
```

```

cache-mode          synchronous
capacity           16G
consistency-group  -
expandable         false
health-indications []
health-state       ok
locality           distributed
operational-status ok
recoverpoint-usage
scsi-release-delay -
service-status     unexported
storage-tier       -
supporting-device  TestDevice
system-id          TestDevice_vol
volume-type        virtual-volume

```

**Table 25 virtual-volume field descriptions (1 of 2)**

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| block count        | The number of blocks in the volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| block size         | The size of a single block, in kilobytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| cache-mode         | Indicates whether the volume use synchronous (write-through) or asynchronous (write-back) IO.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| capacity           | The total number of bytes in the volume. Equals the 'block-size' multiplied by the 'block-count'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| consistency-group  | The name of the consistency group to which this volume belongs, if any.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| expandable         | Whether or not the virtual-volume can be expanded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| health-indications | Indicates the reason(s) for a health-state that is not 'ok'.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| health state       | <p><b>major failure</b> - One or more of the virtual volume's underlying devices is out-of-date, but will never rebuild.</p> <p><b>minor failure</b> - One or more of the virtual volume's underlying devices is out-of-date, but will rebuild.</p> <p><b>non-recoverable error</b> - VPLEX cannot determine the virtual volume's Health state.</p> <p><b>ok</b> - The virtual volume is functioning normally.</p> <p><b>unknown</b> -VPLEX cannot determine the virtual volume's Health state, or the state is invalid.</p>                                                         |
| locality           | <p><b>local</b> - The virtual volume relies completely on storage at its containing cluster.</p> <p><b>remote</b> - The virtual volume is a proxy for a volume whose storage resides at a different cluster. I/O to a remote virtual volume travels between clusters.</p> <p><b>distributed</b> - The virtual volume is the cluster-local representation of a distributed RAID-1. Writes to a distributed volume travels to all the clusters at which it has storage; reads come, if possible, from the local leg.</p>                                                               |
| operational status | <p><b>degraded</b> - The virtual volume may have one or more out-of-date devices that will eventually rebuild.</p> <p><b>error</b> - One or more of the virtual volume's underlying devices is hardware-dead.</p> <p><b>ok</b> - The virtual volume is functioning normally.</p> <p><b>starting</b> -The virtual volume is not yet ready.</p> <p><b>stressed</b> - One or more of the virtual volume's underlying devices is out-of-date and will never rebuild.</p> <p><b>unknown</b> - VPLEX cannot determine the virtual volume's Operational state, or the state is invalid.</p> |
| recoverpoint-usage | The replication role this virtual-volume is being used for by any attached RecoverPoint clusters, if any.<br>Production Source -<br>Local Replica -<br>Remote Replica' -<br>Journal -<br>Repository' -                                                                                                                                                                                                                                                                                                                                                                               |

**Table 25** virtual-volume field descriptions (2 of 2)

| Field              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| scsi-release-delay | A SCSI release delay time in milliseconds. Optimum value is 0 to 2 seconds. Setting a very high value could break the SCSI semantics. If another reserve arrives at this cluster within this time frame, neither release nor reserve will be sent across the WAN.                                                                                                                                                                                                                        |
| service-status     | The service status of a virtual-volume.<br><b>running</b> - I/O is running for the virtual-volume.<br><b>inactive</b> - The virtual-volume is part of an inactive storage-view and is not visible from the host.<br><b>unexported</b> - The virtual-volume is unexported.<br><b>suspended</b> - I/O is suspended for the virtual-volume.<br><b>cluster-unreachable</b> - Cluster is unreachable at this time.<br><b>need-resume</b> - Issue re-attach to resume after link has returned. |
| storage-tier       | The storage-tier for the virtual-volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| supporting-device  | The local, remote, or distributed-device underlying this virtual volume.                                                                                                                                                                                                                                                                                                                                                                                                                 |
| system-id          | The internal system ID for the storage.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| volume-type        | Always virtual-volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

- See also**
- ◆ [virtual-volume destroy on page 507](#)
  - ◆ [virtual-volume expand on page 508](#)
  - ◆ [virtual-volume summary on page 510](#)

---

## virtual-volume destroy

Destroys existing virtual volumes.

**Contexts** All contexts.

**Syntax** `virtual-volume destroy`  
`[-v|--virtual-volumes] context path,context path...`  
`[-f|--force]`

**Arguments** **Required arguments**  
`[-v|--virtual-volumes] context path, context path...` - List of one or more virtual volume(s) to destroy. Entries must be separated by commas. The specified virtual volume(s) must not be exported to hosts.

**Optional arguments**  
`[-f|--force]` - Forces the destruction of the virtual volumes without asking for confirmation. Allows this command to be run from non-interactive scripts.

**Description** Deletes the virtual volume and leaves the underlying structure intact. The data on the volume is no longer accessible.

Only unexported virtual volumes can be deleted. To delete an exported virtual volume, first remove the volume from the storage view.

**Example** `Vplexcli:/clusters/cluster-1> virtual-volume destroy -v`  
`was_1_leg_r1_vol/`

WARNING: The following items will be destroyed:

```
Context
-----
/clusters/cluster-1/virtual-volumes/was_1_leg_r1_vol
Do you wish to proceed? (Yes/No) y
```

**See also**

- ◆ [virtual-volume create on page 504](#)
- ◆ [virtual-volume summary on page 510](#)

## virtual-volume expand

Non-disruptively increases the capacity of an existing virtual volume.

**Contexts** All contexts.

In clusters/*cluster*/virtual-volumes/*virtual-volume* context, command is **expand**.

**Syntax**

```
virtual-volume expand
  [-v|--virtual-volume] context path
  [-e|--extent] extent
  [-f|--force]
```

**Arguments** **Required arguments**

**[-v|--virtual-volume] context path** - \* The virtual volume to expand. The virtual-volume must be expandable, and have a geometry of RAID 1, RAID C, or RAID 0.

**[-e|--extent] context path** - \* The target local device or extent to add to the virtual-volume. The local device or extent must not have a virtual-volume on top of it.

**Optional arguments**

**[-f|--force]** - Forces the expansion of a virtual volume built on a RAID-1 device, even if the target local-device is not of type RAID-1 and/or is not as redundant as the device supporting the virtual volume.

\* - argument is positional.

**Description** Adds the specified extent or device to the virtual volume.

Use the **ll** command in virtual volumes context to display whether the virtual volume is expandable.

RAID C, RAID 1, and RAID 0 devices can be expanded:

- ◆ **RAID C and RAID 0** devices - When expanding a RAID-0 or RAID-C volume, select only one local device or extent to be added.
- ◆ **RAID 1** devices - When expanding a RAID 1 device, select one or multiple volumes to be added.

RAID 1 devices can have up to 8 legs.

Do not expand a virtual volume if either the volume to be expanded or the target to be added is migrating. Wait until the migration is complete.

---

**Note:** Virtual volumes used by RecoverPoint cannot be expanded. To determine if a virtual volume is used by RecoverPoint, type the **ls /clusters/cluster/virtual-volumes/virtual-volume-name** command. If the recoverpoint-usage field is anything other than null (-), the virtual volume cannot be expanded.

---

**Example** In the following example:

- ◆ The **ll clusters/cluster-1/virtual-volumes** command displays the available virtual volumes, and whether the volumes are expandable.
- ◆ The **ll /clusters/cluster-1/storage-elements/extents** command displays available extents.
- ◆ The **cd** command changes the context to that target virtual volume.

- ◆ The **virtual-volume expand --virtual-volume *virtual-volume* --extent *extent*** command adds the specified extent to the specified virtual volume:

```
VPLEXcli: /> ll /clusters/cluster-1/virtual-volumes
/clusters/cluster-1/virtual-volumes:
```

| Name             | Operational Status | Health State | ... | ... | Expandable |
|------------------|--------------------|--------------|-----|-----|------------|
| Raid0_1Ga_11_vol | ok                 | ok           | ... | ... | true       |
| RaidC_1Gb_11_vol | ok                 | ok           | ... | ... | true       |
| Raid1_1Gc_11_vol | ok                 | ok           | ... | ... | true       |
| Test-Device_vol  | ok                 | ok           | ... | ... | true       |

```
.
```

```
.
```

```
.
```

```
VPLEXcli: /> ll /clusters/cluster-1/storage-elements/extends
```

```
/clusters/cluster-1/storage-elements/extends:
```

| Name                       | StorageVolume     | Capacity | Use  |
|----------------------------|-------------------|----------|------|
| extent_Symm1554Tdev_061D_1 | Symm1554Tdev_061D | 100G     | used |
| extent_Symm1554Tdev_0624_1 | Symm1554Tdev_0624 | 100G     | used |
| extent_Symm1554Tdev_0625_1 | Symm1554Tdev_0625 | 100G     | used |
| extent_Symm1554_0690_1     | Symm1554_0690     | 8.43G    | used |
| extent_Symm1554_0691_1     | Symm1554_0691     | 8.43G    | used |
| extent_Symm1554_0692_1     | Symm1554_0692     | 8.43G    | used |

```
.
```

```
.
```

```
.
```

```
VPLEXcli: /> cd /clusters/cluster-1/virtual-volumes/Test-Device_vol
```

```
VPLEXcli:/clusters/cluster-1/virtual-volumes/Test-Device_vol> expand --virtual-volume
Test-Device_vol --extent ext_Symm1254_7BF_1
```

- See also**
- ◆ [virtual-volume create on page 504](#)
  - ◆ [virtual-volume destroy on page 507](#)

## virtual-volume summary

Displays a summary for all virtual volumes.

**Contexts** All contexts.

In `/clusters/cluster-n/virtual-volumes` context, command is **summary**.

**Syntax** `virtual-volume summary`  
`[-c|--clusters] cluster,cluster`

**Arguments** **Required arguments**

None.

**Optional arguments**

`[-c|--clusters] cluster,cluster...` - List of one or more names of clusters. Display information for only the specified cluster(s). Entries must be separated by commas.

**Description** Displays a list of any devices with a health-state or operational-status other than 'ok'.

Displays a summary including devices per locality (distributed versus local), cache-mode (synchronous versus asynchronous), and total capacity for the cluster.

If the `--clusters` argument is not specified and the command is executed at or below `/clusters/cluster` context, information is displayed for the current cluster.

Otherwise, virtual volumes of all clusters are summarized.

**Example** In the following example, cluster 1 has no unhealthy devices (only the summary is displayed), and cluster 2 has one unhealthy device, (the unhealthy device is listed before the summary):

```
Vplexcli:/> virtual-volume summary
Virtual-volume health summary (cluster-1):

    Total 30 virtual-volumes, 0 unhealthy.
Locality summary:
    distributed : 25 virtual-volumes.
    local      : 5 virtual-volumes.
Cache-mode summary:
    synchronous : 30 virtual-volumes.
Total virtual-volume capacity is 2.15T.

Virtual-volume health summary (cluster-2):

volume name                               health state      operational status
-----
sles10_1_esx1_a0_extent_Symm0487_389_1_vol critical-failure  error

    Total 10 volumes, 1 unhealthy.
Locality summary:
    local      : 10 volumes.
Cache-mode summary:
    synchronous : 10 volumes.
    Total virtual-volume capacity is 20.1G.
```

**Table 26 virtual-volume summary field descriptions**

| Field                                                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Virtual-volume health summary (displayed only for unhealthy volumes)</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| volume name                                                                 | Name of the virtual volume.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| health state                                                                | <p><b>major failure</b> - One or more of the virtual volume's underlying devices is out-of-date, but will never rebuild.</p> <p><b>minor failure</b> - One or more of the virtual volume's underlying devices is out-of-date, but will rebuild.</p> <p><b>non-recoverable error</b> - VPLEX cannot determine the virtual volume's Health state.</p> <p><b>ok</b> - The virtual volume is functioning normally.</p> <p><b>unknown</b> -VPLEX cannot determine the virtual volume's Health state, or the state is invalid.</p>                                                         |
| operational status                                                          | <p><b>degraded</b> - The virtual volume may have one or more out-of-date devices that will eventually rebuild.</p> <p><b>error</b> - One or more of the virtual volume's underlying devices is hardware-dead.</p> <p><b>ok</b> - The virtual volume is functioning normally.</p> <p><b>starting</b> -The virtual volume is not yet ready.</p> <p><b>stressed</b> - One or more of the virtual volume's underlying devices is out-of-date and will never rebuild.</p> <p><b>unknown</b> - VPLEX cannot determine the virtual volume's Operational state, or the state is invalid.</p> |
| <b>Summary</b>                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Total                                                                       | Total number of virtual volumes on the cluster, and number of unhealthy virtual volumes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Locality summary                                                            | <p><b>distributed</b> - Number of distributed virtual volumes on the cluster.</p> <p><b>local</b> - Number of local virtual volumes on the cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Cache-mode summary                                                          | How the system's cache flushes writes to the underlying storage volume. For the current release, only synchronous is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

- See also**
- ◆ [ds summary on page 234](#)
  - ◆ [export port summary on page 248](#)
  - ◆ [export storage-view summary on page 264](#)
  - ◆ [extent summary on page 270](#)
  - ◆ [local-device summary on page 283](#)
  - ◆ [virtual-volume summary on page 510](#)

---

## vpn restart

Restarts the VPN connection between management servers.

**Contexts** All contexts.

**Syntax** `vpn restart`

**Description** Restarts the VPN.

**Example** `VPLexcli:/> vpn restart`

**See also**

- ◆ [vpn status on page 514](#)
- ◆ [vpn start on page 513](#)
- ◆ [vpn stop on page 515](#)

---

## vpn start

Starts the VPN connection between management servers.

**Contexts** All contexts.

**Syntax** `vpn start`

**Description** Start the VPN.

**Example** `Vplexcli:/> vpn start`

- See also**
- ◆ [vpn restart on page 512](#)
  - ◆ [vpn status on page 514](#)
  - ◆ [vpn stop on page 515](#)

## vpn status

Verifies the VPN connection between management servers.

**Contexts** All contexts.

**Syntax** `vpn status`

**Description** Verifies whether the VPN connection between the management servers is operating correctly and checks whether all the local and remote directors can be reached (pinged).

If Cluster Witness is deployed, verifies the VPN connection between the management servers and the Cluster Witness Server.

**Example** Display VPN status (no Cluster Witness):

```
Vplexcli:/> vpn status
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.31.25.27 is reachable
Remote Internal Gateway addresses are reachable
Verifying the VPN status between the management server and the cluster witness server...
Cluster Witness Server at IP Address 128.221.254.3 is not reachable
```

**Example** Display VPN status when Cluster Witness is deployed:

```
Vplexcli:/> vpn status
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.31.25.27 is reachable
Remote Internal Gateway addresses are reachable

Verifying the VPN status between the management server and the cluster witness server...
IPSEC is UP
Cluster Witness Server at IP Address 128.221.254.3 is reachable
```

- See also**
- ◆ [About cluster IP seed and cluster ID on page 434](#)
  - ◆ [vpn restart on page 512](#)
  - ◆ [vpn start on page 513](#)
  - ◆ [vpn stop on page 515](#)

---

## vpn stop

Stops the VPN connection between management servers.

**Contexts** All contexts.

**Syntax** `vpn stop`

**Description** Stops the VPN connection between management servers.

**Example** `Vplexcli: /> vpn stop`

- See also**
- ◆ [vpn restart on page 512](#)
  - ◆ [vpn status on page 514](#)
  - ◆ [vpn start on page 513](#)



VPLEX CLI commands, grouped by the following general topics:

- ◆ Storage (arrays, devices, ptov, storage-volumes, virtual volumes)
- ◆ Storage views
- ◆ Data migration
- ◆ System management
- ◆ Clusters
- ◆ Directors
- ◆ Meta-volumes
- ◆ Subnets and ports
- ◆ System status: connectivity, summaries, and reports
- ◆ Diagnostics, dumps, healthchecks
- ◆ Monitoring/statistics
- ◆ Security/authentication
- ◆ Consistency groups
- ◆ Cache vaulting
- ◆ Cluster Witness
- ◆ RecoverPoint
- ◆ Configuration/EZ-Setup
- ◆ Upgrade GeoSynchrony
- ◆ Manage CLI workspace
- ◆ Navigate the CLI
- ◆ Miscellaneous

---

**Storage (arrays, devices, ptov, storage-volumes, virtual volumes)**

advadm dismantle  
array claim  
array re-discover  
copy-session create\*  
device attach-mirror  
device collapse  
device detach-mirror  
device resume-link-down  
device resume-link-up  
ds dd create  
ds dd declare-winner  
ds dd destroy  
ds dd remove-all-rules  
ds dd set-log  
ds rule destroy  
ds rule island-containing  
ds rule island-size  
ds rule-set copy  
ds rule-set create  
ds rule-set destroy  
ds rule-set what-if  
ds summary  
extent create  
extent destroy  
extent summary  
local-device create  
local-device destroy  
local-device summary  
ptov describe-be-zoning  
ptov lun-mask-be-storage  
ptov query-be-storage  
ptov suggest-be-zoning-and-masking  
ptov suggest-ports-to-use  
ptov verify-be-storage  
ptov verify-be-zoning  
rebuild set-transfer-size  
rebuild show-transfer-size  
report capacity-arrays  
set  
storage-volume auto-unbanish-interval  
storage-volume claim  
storage-volume claimingwizard  
storage-volume find-array  
storage-volume forget  
storage-volume list-banished  
storage-volume resurrect  
storage-volume summary  
storage-volume unbanish  
storage-volume unclaim  
storage-volume used-by  
virtual-volume create

virtual-volume destroy  
 virtual-volume expand  
 virtual-volume summary

---

## Storage views

export initiator-port discover  
 export initiator-port register  
 export initiator-port register-host  
 export initiator-port unregister  
 export port summary  
 export storage-view addinitiatorport  
 export storage-view addport  
 export storage-view addvirtualvolume  
 export storage-view checkconfig  
 export storage-view create  
 export storage-view destroy  
 export storage-view find  
 export storage-view find-unmapped-volumes  
 export storage-view map  
 export storage-view removeinitiatorport  
 export storage-view removepor  
 export storage-view removevirtualvolume  
 export storage-view show-powerpath-interfaces  
 export storage-view summary

---

## Data migration

batch-migrate cancel  
 batch-migrate check-plan  
 batch-migrate clean  
 batch-migrate commit  
 batch-migrate create-plan  
 batch-migrate pause  
 batch-migrate remove  
 batch-migrate resume  
 batch-migrate start  
 batch-migrate summary  
 dm migration cancel  
 dm migration clean  
 dm migration commit  
 dm migration pause  
 dm migration remove  
 dm migration resume  
 dm migration start

---

## System management

battery-conditioning disable  
 battery-conditioning enable  
 battery-conditioning manual-cycle cancel-request  
 battery-conditioning manual-cycle request  
 battery-conditioning set-schedule  
 battery-conditioning summary

date  
exec  
exit  
log filter create  
log filter destroy  
log filter list  
log source create  
log source destroy  
log source list  
logging-volume add-mirror  
logging-volume create  
logging-volume destroy  
management-server net-service-restart\*  
manifest upgrade  
manifest version  
notifications call-home import-event-modifications  
notifications call-home remove-event-modifications  
notifications call-home view-event-modifications  
notifications snmp-trap create  
notifications snmp-trap destroy  
remote-clusters add-addresses  
remote-clusters clear-addresses  
report capacity-hosts  
schedule add  
schedule list  
schedule modify  
schedule remove  
scheduleSYR add  
scheduleSYR list  
scheduleSYR remove  
script  
set  
sessions  
snmp-agent configure  
snmp-agent start  
snmp-agent status  
snmp-agent stop  
snmp-agent unconfigure  
source  
user add  
user list  
user passwd  
user remove  
user reset  
version  
vpn restart  
vpn start  
vpn status  
vpn stop

---

## Clusters

cluster add  
cluster cacheflush  
cluster configdump  
cluster expel  
cluster forget

cluster shutdown  
cluster status  
cluster summary  
cluster unexpel  
report capacity-clusters

---

## Directors

connect  
connectivity director  
director appcon  
director commission  
director decommission  
director forget  
director passwd  
director shutdown  
director uptime  
disconnect

---

## Meta-volumes

configuration metadata-backup  
configuration show-meta-volume-candidates  
meta-volume attach-mirror  
meta-volume backup  
meta-volume create  
meta-volume destroy  
meta-volume detach-mirror  
meta-volume move  
meta-volume verify-on-disk-consistency

---

## Subnets and ports

export target-port renamewwns  
set topology  
subnet clear  
subnet create  
subnet destroy  
subnet modify

---

## System status: connectivity, summaries, and reports

cluster summary  
connectivity show  
connectivity validate-be  
connectivity validate-wan-com  
connectivity window set  
connectivity window show  
connectivity window stat  
consistency-group summary  
director ping  
director tracepath  
ds summary

event-test  
 export port summary  
 export storage-view checkconfig  
 export storage-view summary  
 extent summary  
 local-device summary  
 ls  
 notifications call-home test  
 rebuild show-transfer-size  
 rebuild status  
 report capacity-arrays  
 report capacity-clusters  
 report capacity-hosts  
 storage-volume summary  
 validate-system-configuration  
 verify fibre-channel-switches  
 virtual-volume summary

---

## Diagnostics, dumps, healthchecks

collect-diagnostics  
 collect-recoverpoint-diagnostics  
 director appdump  
 director appstatus  
 director firmware show-banks  
 director ping  
 getsysinfo  
 health-check  
 sms dump  
 syrcollect

---

## Monitoring/statistics

monitor add-console-sink  
 monitor add-file-sink  
 monitor collect  
 monitor create  
 monitor destroy  
 monitor remove-sink  
 monitor stat-list  
 report aggregate-monitors  
 report create-monitors  
 report poll-monitors

---

## Security/authentication

authentication directory-service configure  
 authentication directory-service map  
 authentication directory-service show  
 authentication directory-service unconfigure  
 authentication directory-service unmap  
 configuration configure-auth-service  
 security create-ca-certificate  
 security create-certificate-subject

security create-host-certificate  
 security delete-ca-certificate  
 security delete-hostcertificate  
 security export-ca-certificate  
 security export-host-certificate  
 security import-ca-certificate  
 security import-host-certificate  
 security renew-all-certificates  
 security ipsec-configure  
 security show-cert-subj

---

## Consistency groups

consistency-group add-virtual-volumes  
 consistency-group choose-winner  
 consistency-group create  
 consistency-group destroy  
 consistency-group list-eligible-volumes  
 consistency-group remove-virtual-volumes  
 consistency-group resolve-conflicting-detach  
 consistency-group resume-after-data-loss-failure  
 consistency-group resume-after-rollback  
 consistency-group resume-at-loser  
 consistency-group set-detach-rule active-cluster-wins  
 consistency-group set-detach-rule no-automatic-winner  
 consistency-group set-detach-rule winner  
 consistency-group summary  
 set

---

## Cache vaulting

vault go  
 vault overrideUnvaultQuorum  
 vault status

---

## Cluster Witness

cluster-witness configure  
 cluster-witness disable  
 cluster-witness enable  
 configuration cw-vpn-configure  
 configuration cw-vpn-reset

---

## RecoverPoint

rp rpa-cluster add  
 rp rpa-cluster remove  
 rp summary  
 rp validate-configuration

---

## Configuration/EZ-Setup

configuration connect-local-directors

configuration complete-system-setup  
configuration connect-local-directors  
configuration connect-remote-directors  
configuration continue-system-setup  
configuration enable-front-end-ports  
configuration event-notices-reports config  
configuration event-notices-reports reset  
configuration get-product-type  
configuration join-clusters  
configuration register-product  
configuration system-reset  
configuration system-setup

configuration cw-vpn-configure  
configuration cw-vpn-reset  
configuration show-meta-volume-candidates

---

## Upgrade GeoSynchrony

ndu pre-config-upgrade  
ndu pre-check  
ndu recover  
ndu rolling-upgrade c4lx-to-sles  
ndu rolling-upgrade ssd-fw  
ndu start  
ndu status

---

## Manage CLI workspace

alias  
capture begin  
capture end  
capture pause  
capture replay  
capture resume  
describe  
dirs  
drill-down  
unalias  
security remove-login-banner  
security set-login-banner

---

## Navigate the CLI

cd  
help  
history  
ls  
popd  
pushd  
tree

---

## Miscellaneous

alias  
chart create



**A**

- add-console-sink, monitor 315
- add-file-sink, monitor 317
- addinitiatorport, export view 250
- add-mirror, logging-volume 292
- addport, export view 251
- addurl, plugin 368
- advirtualvolume, export view 252
- advadm dismantle 26
- aggregate, report monitors 386
- appdump, director 190
- array claim 29
- array re-discover 30
- array used-by 32
- attach-mirror
  - meta-volume 303
- authentication 34, 38, 40, 41
- authentication directory-service 39
- authentication directory-service configure 34
- authentication directory-service map 38
- authentication directory-service show 39
- authentication directory-service unconfigure 40

**B**

- Batch migration
  - perform a migration 48
- batch-migrate
  - check-plan 43
  - pause 50
  - remove 51
  - resume 52
  - start 53
  - summary 55
- batch-migrate cancel 42
- battery-conditioning disable 57
- battery-conditioning enable 59
- battery-conditioning manual-cycle cancel-request 61
- battery-conditioning manual-cycle request 63
- battery-conditioning set-enabled 59
- battery-conditioning set-schedule 65
- battery-conditioning summary 67
- begin, capture 69

**C**

- call-home import-event-modifications 359
- call-home remove-event-modifications 361
- call-home-view-event-modifications 363
- cancel
  - dm migration 207
- capacity-clusters, report 390
- capacity-hosts, report 391
- capture
  - begin 69
  - end 70
  - pause 71
  - replay 72
  - resume 73
- checkconfig, export view 254
- check-plan, batch-migrate 43
- claim, storage-volume 458
- claimingwizard 461
- clean, dm migration 208
- CLI
  - ? wildcard 19
  - \* wildcard 19
  - \*\* wildcard 19
  - accessing help 22
  - context tree 11, 13
  - display available commands 17
  - display command arguments 19
  - executing commands 17
  - Log in 10
  - Log out 10
  - log out 10
  - login 10
  - page output 18
  - positional arguments 21
  - Tab completion 18
  - using commands 17
  - verbose and --help arguments 21
  - viewing command history 22
- CLI commands
  - cd 15
  - globbing 20
  - ls 14
  - popd 13

- positional arguments 21
  - pushd 13
  - tree 15
  - verbose and --help arguments 21
  - wildcards 19
  - CLI context tree 13
  - CLI workspace
    - command history 22
  - cluster
    - capacity report 390
    - summary 93
  - cluster configdump 81
  - cluster IP seed 434
  - cluster status
    - cluster shutdown 90
    - expelled cluster 90
    - healthy clusters 89
    - link outage 89
  - cluster summary
    - cluster expelled 93
    - healthy cluster 93
    - link outage 93
  - cluster-witness
    - configure 96
    - create context 96
    - enable 101
    - VPN connectivity 115
    - vpn prerequisites 115
  - cluster-witness disable 98
  - collect
    - monitor 319
  - collect-diagnostics 106
  - command history, viewing 22
  - commission, director 193
  - commit, dm migration 209
  - configuration complete-geo-setup 109
  - configuration configure-auth-service 110
  - configuration connect-local-directors 112
  - configuration cws-vpn-configure 115
  - configuration enable-front-end-ports 120
  - configuration event-notices-reports-config 121
  - configuration event-notices-reports-reset 124
  - configuration metadata-backup 128
  - configuration register-product 130
  - configuration show-meta-volume-candidates 132
  - configuration sync-time 133
  - configuration system-reset 134
  - configuration system-setup 113, 114, 135
  - configuration, validate 492
  - connect 137
  - connectivity
    - show 141
    - validate-be 142
    - window set 147
    - window show 149
    - window stat 150
  - connectivity validate-wan-com 145
  - consistency-group choose-winner 153
  - consistency-group create 156
  - consistency-group destroy 159
  - consistency-group list-eligible-volumes 160
  - consistency-group remove-virtual-volumes 161
  - consistency-group resolve-conflicting-detach 163
  - consistency-group resume-after-data-loss-failure 166
  - consistency-group resume-after-rollback 169
  - consistency-group resume-at-loser 172
  - consistency-group summary 178
  - console sessions 443
  - copy rule-set 230
  - create
    - distributed-device 219
    - export view 255
    - local-device 279
    - log filter 285
    - log source 289
    - logging-volume 293
    - meta-volume 306
    - monitor 320
  - create extent 267
  - create, snmp-trap 366
  - create-ca-certificate
    - security 417
  - create-cert-subj
    - security 421
  - create-host-certificate
    - security 423
  - create-monitors, report 392
- ## D
- declare-winner, distributed-device 222
  - destroy
    - distributed-device 223
    - export view 256
    - extent 269
    - local-device 282
    - log filter 287
    - log source 290
    - logging-volume 295
    - meta-volume 310
    - monitor 323
    - rule 227
    - rule-set 232
    - volume 507
  - destroy, snmp-trap 367
  - Detach
    - mirror leg 184
  - detach-mirror, device 184
  - device
    - create local 279
    - destroy local 282
    - detach-mirror 184
    - resume-link-down 186
    - resume-link-up 188
  - director
    - appdump 190
    - commission 193
    - passwd 199
    - ping 200
    - shutdown 201
  - director appcon 189

- director fc-port-stats 195
  - director firmware show-banks 197
  - director tracepath 202
  - directory-serevice unconfigure 40
  - directory-service map 38
  - directory-service unmap 41
  - dirs 205
  - disconnect 206
  - discovering
    - initiator ports 243
  - dismantle, advadm 26
  - Display
    - default values 444
    - multiple objects using globs 297
    - object's attributes 296
    - object's attributes and contexts 297
    - operational status 297
    - writable attributes 444
  - distributed-device
    - create 219
    - declare-winner 222
    - destroy 223
    - set-log 225
  - distributed-storage
    - dd create 219
    - dd declare-winner 222
    - dd destroy 223
    - dd set-log 225
    - rule destroy 227
    - rule island-containing 228
    - rule island-size 229
    - rule-set copy 230
    - rule-set destroy 232
  - dm
    - migration cancel 207
    - migration clean 208
    - migration commit 209
    - migration remove 211
    - migration resume 212
    - migration start 213
  - dr
    - create 219
    - declare-winner 222
    - destroy 223
    - set-log 225
  - drill-down 217
  - ds
    - dd create 219
    - dd declare-winner 222
    - dd destroy 223
    - dd set-log 225
    - rule destroy 227
    - rule island-containing 228
    - rule island-size 229
    - rule-set copy 230
    - rule-set destroy 232
- E**
- end, capture 70
  - event-test 239
  - exec 241
  - executing CLI commands 17
  - exit 242
  - export
    - initiator-port register 244
    - initiator-port register-host 246
    - initiator-port unregister 247
    - port summary 248
    - target-port renamewwns 266
    - view addinitiatorport 250
    - view addport 251
    - view addvirtualvolume 252
    - view checkconfig 254
    - view create 255
    - view destroy 256
    - view find 257
    - view find-unmapped-volumes 258
    - view map 259
    - view removeinitiatorport 260
    - view removeport 261
    - view removevirtualvolume 262
    - view show-powerpath-interfaces 263
    - view summary 264
  - export initiator-port discovery 243
  - extent
    - create 267
    - destroy 269
    - summary 270
    - unclaim 272
  - Extent migrations 214
- F**
- file rotation 321
  - find, export view 257
  - find-array, storage-volume 464
  - find-unmapped-volumes, export view 258
  - forget
    - storage-volume 465
- G**
- gateway IP address
    - set 299
  - getsysinfo 272
- H**
- help 276
  - help command
    - CLI 22
  - help, accessing in CLI 22
  - history 278
  - host, capacity report 391
- I**
- import-event-modifications 359
  - initiator ports
    - discovering 243
  - initiator-port

- register 244
- register-host 246
- unregister 247
- IP address
  - set 299
- IP seed 434
- ipsec-configure
  - security 433
- island-containing rule 228
- island-size rule 229

**J**

- Jython 416

**L**

- list 296
  - log filter 288
  - log source 291
- listurl, plugin 369
- local-device
  - create 279
  - destroy 282
  - summary 283
- log
  - filter create 285
  - filter destroy 287
  - filter list 288
  - source create 289
  - source destroy 290
  - source list 291
- Log in to CLI 10
- Log out from CLI 10
- log source create 289
- log source destroy 290
- logging-volume
  - add-mirror 292
  - create 293
  - destroy 295
- ls 296

**M**

- management-server
  - set-ip 299
- management-server set-ip 299
- manifest
  - upgrade 301
  - version 302
- map, export view 259
- metadata
  - backup
    - initial 128
    - modify schedule 129
  - modify backup schedule 129
  - schedule backup 128
- metadata backup
  - display backup volumes 129
- meta-volume
  - attach-mirror 303
  - create 306

- destroy 310
- move 312
- migration
  - cancel 207
  - clean 208
  - commit 209
  - remove 211
  - resume 212
  - start 213
- migrations
  - device 214
  - extent 214
  - prerequisites 214
- mirror leg
  - detach 184
- monitor
  - add-console-sink 315
  - add-file-sink 317
  - collect 319
  - create 320
  - destroy 323
  - remove-sink 324
  - stat-list 325
- monitors, aggregate 386
- move, meta-volume 312

**N**

- Names 20
  - VPLEX components 20
- NDU
  - Before you begin 351
  - display status 357
  - force-geo option 350
  - pre-checks and skip options 351
  - pre-config-upgrade 327
    - procedure 327
  - pre-ndu validation 330
  - recover from failed ndu 335
  - rolling upgrade (directors) 337
  - rolling upgrade SSD firmware 341
  - upgrade unconfigured Geo 327
- ndu pre-check 330
- ndu pre-config-upgrade 327
- ndu recover 335
- ndu rolling-upgrade c41-to-sles 337
- ndu rolling-upgrade c41x-to-sles 337
- ndu rolling-upgrade ssd-fw 341
- ndu start 348
- ndu status 357
- net-mask
  - set 299
- Network
  - set IP address 299
- notifications call-home importconfig 359
- notifications call-home test 365
- notifications call-home view-config 363

**P**

- passwd, director 199

- pause
  - batch-migrate 50
  - capture 71
- Performance monitoring
  - file rotation 321
- physical-to-virtual
  - list zones 372
  - mask LUNs 373
  - query storage 374
  - suggest ports 377
  - suggest zoning 376
  - verify BE connectivity 378
  - verify zoning 379
- ping 200
- plugin
  - addurl 368
  - listurl 369
  - register 370
- poll-monitors, report 395
- popd 371
- port, export summary 248
- ptov describe-be-zoning 372
- ptov lun-mask-be-storage 373
- ptov query-be-storage 374
- ptov suggest-be-zoning-and-masking 376
- ptov suggest-ports-to-use 377
- ptov verify-be-storage 378
- ptov verify-be-zoning 379
- pushd 380

## R

- rebuild
  - setlimit 381
  - showlimit 382
  - status 383
- RecoverPoint Appliance
  - remove 402
- Re-discover an array 30
- register
  - initiator-port 244
  - plugin 370
- register-host, export initiator-port 246
- remote-clusters add-addresses 384
- remote-clusters clear-addresses 385
- remove
  - batch-migrate 51
  - dm migration 211
- removeinitiatorport, export view 260
- remove-modifications 361
- removeport, export view 261
- remove-sink, monitor 324
- removevirtualvolume, export view 262
- renamewwns, export target-port 266
- Renew security certificates 436
- replay, capture 72
- report
  - aggregate-monitors 386
  - capacity-clusters 390
  - capacity-hosts 391
  - create-monitors 392

- poll-monitors 395
- restart VPN 512
- resume
  - batch-migrate 52
  - capture 73
  - dm migration 212
- resume-link-down, device 186
- resume-link-up, device 188
- resurrect, storage-volume 467
- rp import-certificate 396
- rp rpa add 397
- rp rpa remove 402
- rp rpa summary 403
- rp validate configuration 405
- rule
  - destroy 227
  - island-containing 228
  - island-size 229
- rule-set
  - copy 230
  - destroy 232

## S

- schedule add 409
- schedule modify 411
- schedule remove 412
- script, command 416
- security
  - create-ca-certificate 417
  - create-cert-subj 421
  - create-host-certificate 423
  - ipsec-configure 433
  - renew-all-certificates 436
  - show-cert-subj 442
- security delete-ca-certificate 425
- security delete-host-certificate 426
- security remove-login-banner 435
- security renew-all-certificates 436
- security set-login-banner 441
- set
  - connectivity window 147
- set topology 447
- set, command 444
- set-ip
  - management-server 299
- setlimit, rebuild 381
- set-log, distributed-device 225
- show 39, 443
  - connectivity 141
  - connectivity window 149
- show-cert-subj
  - security 442
- showlimit, rebuild 382
- show-powerpath-interfaces, export view 263
- shutdown
  - director 201
- sink
  - add-console 315
  - add-file 317
  - remove 324

- sms dump 448
- snmp-agent configure 450
- snmp-agent start 451
- snmp-agent status 452
- snmp-agent stop 453
- snmp-agent unconfigure 454
- snmp-trap, create 366
- snmp-trap, destroy 367
- source 455
- start
  - batch-migrate 53
  - dm migration 213
- start VPN 513
- stat, connectivity window 150
- stat-list, monitor 325
- status
  - rebuild 383
- status VPN 514
- stop VPN 515
- storage tier IDs 504
- storage tiers 504
- storage-volume
  - claim 458
  - claimingwizard 461
  - find-array 464
  - forget 465
  - resurrect 467
  - summary 469
  - unclaim 474
  - used-by 475
- subnet clear 476, 480
- subnet create 478
- subnet destroy 480
- subnet modify 481
- summary
  - batch-migrate 55
  - cluster 93
  - export port 248
  - export view 264
  - extent 270
  - local-device 283
  - storage-volume 469
  - volume 510
- system, validate configuration 492

**T**

- target-port, renamewwns 266
- test, call-home notifications 365
- Thin provisioning 458
- tree 484

**U**

- unalias 485
- unclaim extent 272
- unclaim, storage-volume 474
- unregister, initiator-port 247
- upgrade
  - manifest 301
- used-by, storage-volume 475

- user reset 491

**V**

- validate-be, connectivity 142
- validate-system-configuration 492
- vault go 493
- vault overrideUnvaultQuorum 494
- verify fibre-channel-switches 500
- version 501
  - manifest 302
- view
  - addinitiatorport 250
  - addport 251
  - addvirtualvolume 252
  - checkconfig 254
  - create 255
  - destroy 256
  - find 257
  - find-unmapped-volumes 258
  - map 259
  - removeinitiatorport 260
  - removeport 261
  - removevirtualvolume 262
  - show-powerpath-interfaces 263
  - summary 264
- view-event-modifications 363
- virtual-volume expand 508
- volume
  - destroy 507
  - summary 510
- vpn restart 512
- vpn start 513
- vpn status 514
- vpn stop 515

**W**

- WAN COM
  - verify connectivity 145
- wizard, claiming 461