

Dell EMC VPLEX™ GeoSynchrony

Version Release 6.2

CLI Reference Guide

Rev 01

Copyright © 2019 Dell Inc. or its subsidiaries. All rights reserved.

Dell 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.” DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Tables		11
	Preface	13
	New commands in this release	17
	Deprecated CLI commands	19
Chapter 1	Using the VPLEX CLI	21
	Logging in to the CLI.....	22
	Password Policies.....	23
	Valid Password Characters.....	24
	Logging out from the CLI.....	24
	CLI context tree.....	24
	Context tree searching.....	25
	Navigate the CLI context tree.....	26
	pushd and popd commands.....	27
	Where am I in the context tree?.....	27
	Using CLI commands.....	28
	Page output.....	29
	Tab completion.....	29
	Wildcards.....	31
	Names.....	32
	Specifying addresses.....	32
	Command globbing.....	33
	Positional command arguments.....	33
	Search command history.....	34
	View command history.....	34
	Get help.....	35
Chapter 2	Commands	37
	advadm dismantle.....	44
	alias.....	45
	amp register.....	47
	amp unregister.....	49
	array claim.....	50
	array forget.....	51
	array re-discover.....	52
	array used-by.....	54
	authentication directory-service configure.....	55
	authentication directory-service map.....	59
	authentication directory-service show.....	60
	authentication directory-service unconfigure.....	62
	authentication directory-service unmap.....	62
	back-end degraded list.....	64
	back-end degraded recover.....	65
	batch-migrate cancel.....	66

batch-migrate check-plan.....	67
batch-migrate clean.....	69
batch-migrate commit.....	70
batch-migrate create-plan.....	71
batch-migrate pause.....	73
batch-migrate remove.....	74
batch-migrate resume.....	75
batch-migrate start.....	76
batch-migrate summary.....	78
battery-conditioning disable.....	81
battery-conditioning enable.....	83
battery-conditioning manual-cycle cancel-request.....	85
battery-conditioning manual-cycle request.....	86
battery-conditioning set-schedule.....	87
battery-conditioning summary.....	90
cache-invalidate.....	92
cache-invalidate-status.....	95
capture begin.....	98
capture end.....	99
capture pause.....	100
capture replay.....	100
capture resume.....	101
cd.....	102
chart create.....	103
cluster add.....	105
cluster cacheflush.....	106
cluster configdump.....	107
cluster expel.....	109
cluster forget.....	110
cluster restart-local-cluster.....	111
cluster show-remote-devices.....	112
cluster shutdown.....	114
cluster status.....	116
cluster stop-local-cluster.....	119
cluster summary.....	119
cluster unexpel.....	123
cluster-witness configure.....	124
cluster-witness disable.....	126
cluster-witness enable.....	128
collect-diagnostics.....	133
configuration complete-system-setup.....	135
configuration configure-auth-service.....	136
configuration connect-local-directors.....	137
configuration connect-remote-directors.....	138
configuration continue-system-setup.....	139
configuration cw-vpn-configure.....	139
configuration cw-change-password.....	141
configuration cw-vpn-reset.....	141
configuration enable-front-end-ports.....	143
configuration event-notices-reports config.....	144
configuration event-notices-reports reset.....	144
configuration event-notices-reports-show.....	145
configuration flashdir-backup disable.....	146
configuration flashdir-backup enable.....	147
configuration get-product-type.....	147
configuration join-clusters.....	148

configuration metadata-backup.....	149
configuration register-product.....	151
configuration remote-clusters add-addresses.....	152
configuration remote-clusters clear-addresses.....	154
configuration show-meta-volume-candidates.....	155
configuration subnet clear.....	156
configuration subnet remote-subnet add.....	158
configuration subnet remote-subnet remove.....	159
configuration sync-time.....	161
configuration sync-time-clear.....	162
configuration sync-time-show.....	163
configuration system-reset.....	164
configuration system-setup.....	165
configuration upgrade-meta-slot-count.....	166
connect.....	167
connectivity director.....	169
connectivity list all.....	170
connectivity list directors.....	171
connectivity list initiators.....	172
connectivity list storage-volumes.....	173
connectivity show.....	174
connectivity validate-be.....	174
connectivity validate-local-com.....	177
connectivity validate-wan-com.....	178
consistency-group add-virtual-volumes.....	179
consistency-group choose-winner.....	180
consistency-group convert-to-local.....	183
consistency-group create.....	184
consistency-group destroy.....	186
consistency-group list-eligible-virtual-volumes.....	187
consistency-group remove-virtual-volumes.....	188
consistency-group resolve-conflicting-detach.....	190
consistency-group resume-at-loser.....	192
consistency-group set-detach-rule no-automatic-winner.....	194
consistency-group set-detach-rule winner.....	195
consistency-group summary.....	197
date.....	197
describe.....	198
device attach-mirror.....	199
device collapse.....	201
device detach-mirror.....	202
device mirror-isolation auto-unisolation disable.....	206
device mirror-isolation auto-unisolation enable.....	207
device mirror-isolation disable.....	209
device mirror-isolation enable.....	211
device mirror-isolation show.....	213
device resume-link-down.....	215
device resume-link-up.....	217
device resurrect-dead-storage-volumes.....	218
director appcon.....	219
director appdump.....	220
director appstatus.....	222
director commission.....	223
director decommission.....	224
director fc-port-stats.....	224
director firmware show-banks.....	226

director forget.....	227
director passwd.....	228
director ping.....	229
director shutdown.....	230
director tracepath.....	231
director uptime.....	232
dirs.....	233
disconnect.....	233
dm migration cancel.....	234
dm migration clean.....	235
dm migration commit.....	236
dm migration pause.....	237
dm migration remove.....	238
dm migration resume.....	239
dm migration start.....	240
drill-down.....	242
ds dd convert-to-local.....	244
ds dd create.....	245
ds dd declare-winner.....	247
ds dd destroy.....	248
ds dd remove-all-rules.....	249
ds dd set-log.....	250
ds rule destroy.....	252
ds rule island-containing.....	252
ds rule-set copy.....	254
ds rule-set create.....	255
ds rule-set destroy.....	255
ds rule-set what-if.....	257
ds summary.....	258
esrs import-certificate.....	265
esrs register.....	266
esrs status.....	267
esrs un-register.....	268
event-test.....	269
exec.....	270
exit.....	270
export initiator-port discovery.....	271
export initiator-port register.....	272
export initiator-port register-host.....	273
export initiator-port show-logins.....	274
export initiator-port unregister.....	276
export port summary.....	277
export storage-view addinitiatorport.....	279
export storage-view addport.....	280
export storage-view addvirtualvolume.....	280
export storage-view checkconfig.....	283
export storage-view create.....	284
export storage-view destroy.....	285
export storage-view find.....	286
export storage-view find-unmapped-volumes.....	287
export storage-view map.....	288
export storage-view removeinitiatorport.....	289
export storage-view removeport.....	290
export storage-view removevirtualvolume.....	291
export storage-view show-powerpath-interfaces.....	292
export storage-view summary.....	292

export target-port renamewwns.....	293
extent create.....	294
extent destroy.....	296
extent summary.....	297
find.....	299
front-end-performance-stats start.....	300
front-end-performance-stats status.....	300
front-end-performance-stats stop.....	301
getsysinfo.....	302
health-check.....	303
help.....	309
history.....	310
iscsi chap back-end add-credentials.....	311
iscsi chap back-end disable.....	312
iscsi chap back-end enable.....	313
iscsi chap back-end list-credentials.....	314
iscsi chap back-end remove-credentials.....	315
iscsi chap back-end remove-default-credential.....	316
iscsi chap back-end set-default-credential.....	317
iscsi chap front-end add-credentials.....	318
iscsi chap front-end disable.....	320
iscsi chap front-end enable.....	321
iscsi chap front-end list-credentials.....	322
iscsi chap front-end remove-credentials.....	323
iscsi chap front-end remove-default-credential.....	324
iscsi chap front-end set-default-credential.....	325
iscsi check-febe-connectivity.....	326
iscsi isns add.....	327
iscsi isns list.....	328
iscsi isns remove.....	330
iscsi sendtargets add.....	331
iscsi sendtargets list.....	332
iscsi sendtargets rediscover.....	334
iscsi sendtargets remove.....	335
iscsi targets list.....	336
iscsi targets logout.....	337
license install.....	338
license reset.....	339
license show.....	340
local-device create.....	341
local-device destroy.....	344
local-device summary.....	345
log filter create.....	346
log filter destroy.....	348
log filter list.....	348
log source create.....	349
log source destroy.....	350
log source list.....	351
logging-volume add-mirror.....	351
logging-volume create.....	352
logging-volume detach-mirror.....	357
logging-volume destroy.....	359
logical-unit forget.....	360
ls.....	361
management-server set-ip.....	364
manifest upgrade.....	366

manifest version.....	366
meta-volume attach-mirror.....	367
meta-volume backup.....	368
meta-volume create.....	369
meta-volume destroy.....	372
meta-volume detach-mirror.....	373
meta-volume move.....	374
meta-volume verify-on-disk-consistency.....	375
monitor add-console-sink.....	376
monitor add-file-sink.....	378
monitor collect.....	379
monitor create.....	380
monitor destroy.....	383
monitor get-stats.....	384
monitor remove-sink.....	385
monitor stat-list.....	385
ndu pre-check.....	386
ndu pre-config-upgrade.....	387
ndu recover.....	388
ndu rolling-upgrade ssd-fw.....	388
ndu start.....	391
ndu status.....	394
notifications call-home import-event-modifications.....	395
notifications call-home remove-event-modifications.....	395
notifications call-home view-event-modifications.....	396
notifications call-home test.....	397
notifications job cancel.....	398
notifications job delete.....	398
notifications job resubmit.....	400
notifications snmp-trap create.....	401
notifications snmp-trap destroy.....	402
password-policy reset.....	402
password-policy set.....	403
plugin addurl.....	405
plugin listurl.....	406
plugin register.....	407
popd.....	407
pushd.....	408
rebuild set-transfer-size.....	409
rebuild show-transfer-size.....	410
rebuild status.....	411
report aggregate-monitors.....	412
report capacity-arrays.....	413
report capacity-clusters.....	416
report capacity-hosts.....	417
report create-monitors.....	419
report poll-monitors.....	421
rm.....	422
rp import-certificate.....	423
rp rpa-cluster add.....	424
rp rpa-cluster remove.....	431
rp summary.....	432
rp validate-configuration.....	434
schedule add.....	437
schedule list.....	438
schedule modify.....	438

schedule remove.....	439
scheduleSYR add.....	440
scheduleSYR list.....	441
scheduleSYR remove.....	442
script.....	442
security configure-mmcs-users.....	443
security create-ca-certificate.....	444
security create-certificate-subject.....	446
security create-host-certificate.....	448
security configure-certificates.....	450
security delete-ca-certificate.....	454
security delete-host-certificate.....	455
security export-ca-certificate.....	456
security export-host-certificate.....	457
security import-ca-certificate.....	458
security import-host-certificate.....	459
security ipsec-configure.....	461
security list-certificates.....	463
security remove-login-banner.....	464
security renew-all-certificates.....	465
security set-login-banner.....	469
security show-cert-subj.....	470
security web-configure.....	471
sessions.....	471
set.....	472
set topology.....	477
show-use-hierarchy.....	478
sms dump.....	479
snmp-agent configure.....	481
snmp-agent start.....	482
snmp-agent status.....	483
snmp-agent stop.....	483
snmp-agent unconfigure.....	484
source.....	484
storage-tool dismantle.....	485
storage-tool compose.....	486
storage-volume auto-unbanish-interval.....	489
storage-volume claim.....	490
storage-volume claimingwizard.....	493
storage-volume find-array.....	496
storage-volume forget.....	498
storage-volume list-banished.....	499
storage-volume list-thin-capable.....	500
storage-volume resurrect.....	501
storage-volume summary.....	503
storage-volume unbanish.....	507
storage-volume unclaim.....	508
storage-volume used-by.....	510
syrcollect.....	511
tree.....	512
unalias.....	513
user add.....	514
user event-server add-user.....	515
user event-server change-password.....	516
user list.....	517
user passwd.....	518

user remove..... 519

user reset..... 520

validate-system-configuration..... 521

vault go..... 522

vault overrideUnvaultQuorum..... 523

vault status..... 525

verify fibre-channel-switches..... 529

version..... 529

virtual-volume create..... 533

virtual-volume destroy..... 540

virtual-volume expand..... 541

virtual-volume list-thin..... 545

virtual-volume provision..... 546

virtual-volume re-initialize..... 548

virtual-volume set-thin-enabled..... 548

virtual-volume summary..... 549

vpn restart..... 553

vpn start..... 553

vpn status..... 554

vpn stop..... 555

wait..... 555

webserver..... 556

Index

557


TABLES

1	Typographical conventions.....	14
2	Deprecated options to commands.....	19
3	Default password policies.....	23
4	authentication directory-service show field descriptions.....	61
5	batch migration summary field descriptions.....	79
6	battery conditioning field descriptions.....	88
7	battery conditioning summary field descriptions	90
8	Important cache invalidate status fields.....	96
9	cluster status field descriptions.....	116
10	cluster summary field descriptions.....	120
11	cluster witness display fields.....	129
12	Connection Types.....	145
13	Supported configuration parameters.....	214
14	director firmware show-banks field descriptions	227
15	ds summary field descriptions	258
16	extent summary field descriptions	298
17	getsysinfo field descriptions.....	302
18	local device summary field descriptions	345
19	logging volume display fields.....	353
20	recoverpoint display fields.....	425
21	rp summary display fields	432
22	Certificate parameters.....	463
23	Create hints files for storage-volume naming	494
24	storage-volume summary field descriptions.....	503
25	Vault state field descriptions.....	526
26	Software components.....	530
27	virtual-volume field descriptions.....	534
28	virtual-volume summary field descriptions.....	550

Preface

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

 **Note:** This document was accurate at publication time. Go to Dell EMC Online Support (<https://www.dell.com/support>) to ensure that you are using the latest version of this document.

Purpose

This document is part of the VPLEX documentation set, and describes the VPLEX features and use cases, configuration options, VPLEX software and its upgrade, and the hardware overview.

Audience


This guide is intended for use by customers who wish to understand the software and hardware features of VPLEX, the use cases of VPLEX, product offerings, and the configuration options.


Related documents (available on Dell EMC Online Support) include:


- *VPLEX Release Notes for GeoSynchrony Releases*
- *VPLEX Product Guide*
- *VPLEX Hardware Environment Setup Guide*
- *VPLEX Configuration Worksheet*
- *VPLEX Configuration Guide*
- *VPLEX Security Configuration Guide*
- *VPLEX CLI Reference Guide*
- *VPLEX Administration Guide*
- Unisphere for VPLEX Help
- *VPLEX Element Manager API Guide*
- *VPLEX Open-Source Licenses*
- *VPLEX GPL3 Open-Source Licenses*
- Procedures provided through the SolVe Desktop
- Dell EMC Host Connectivity Guides
- *Dell EMC VPLEX Hardware Installation Guide*
- Various best practices technical notes available on Dell EMC Online Support

Special notice conventions used in this document


Dell EMC uses the following conventions for special notices:

 **DANGER** Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

 **WARNING** Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

 **CAUTION** Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

 **NOTICE** Addresses practices not related to personal injury.

 **Note:** Presents information that is important, but not hazard-related.

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Table 1 Typographical conventions

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none"> • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
<i>Monospace italic</i>	Used for variables
Monospace bold	Used for user input
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to Dell EMC Online Support at <https://www.dell.com/support>.

Technical support

Go to Dell EMC Online Support and click Support. You will see several options for contacting Dell EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Online communities

Visit Dell EMC Community Network (DECN) at <https://www.dell.com/community/Dell-Community/ct-p/English> for peer contacts, conversations, and content on product support and

solutions. Interactively engage online with customers, partners, and certified professionals for all Dell EMC products.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

New commands in this release

The following commands and contexts have been added, removed, or changed in this release.

New commands in this release

The following commands have been added to this release:

- `back-end degraded list`
- `back-end degraded recover`
- `esrs import-certificate`
- `esrs register`
- `esrs status`
- `esrs un-register`

Changed commands

For this release, there are no changes in the commands.

Removed commands

Starting with this release, `connectivity window set` and `connectivity window show` commands are removed.

New or changed contexts

Starting with this release, the cluster configuration limits are updated for Health-check limits on a metro system .

New commands in this release

Deprecated CLI commands

To increase the overall scale and performance of VPLEX and to reduce the product complexity, commands or contexts might be deprecated and contexts might be changed. These options and arguments will be removed from the CLI in the future releases of VPLEX.

The deprecated options and arguments and the workaround to mitigate the deprecation are as follows:

Table 2 Deprecated options to commands

Command	Deprecated options/arguments	Description	Workaround
extent create	optional argument: [-n --num-extents]	Specifies the number of extents to be created on a specific storage volume. When this argument is not used, VPLEX creates a single extent for specific storage volume.	Use array-native slicing capabilities.
	optional argument: [-o --block-offset]	Specifies the block-offset on the underlying storage volume on which the extent is created. If not specified, the block-offset is determined automatically.	None needed.
local-device create	The <code>raid-0</code> value of the optional argument [-g --geometry]	Specifies the creation of a local RAID 0 device.	Use array-native striping capabilities.
	Optional argument: [-d --stripe-depth]	Specifies the stripe depth for a RAID 0 device.	None needed.
local-device create	The <code>raid-c</code> value of the optional argument [-g --geometry].	Specifies the creation of a local RAID C device.	Use array-native striping capabilities.
logging-volume create	The <code>raid-0</code> value of the optional argument [-g --geometry]	Specifies the creation of a logging volumes using RAID 0/striping, internally.	Use array-native striping capabilities.
	Optional argument: [-d --stripe-depth]	Specifies the stripe depth for a RAID 0 device.	None needed.
storage-tool compose	The <code>raid-0</code> value of the optional argument [-g --geometry] .	Specifies the creation of a local RAID 0 device.	Use array-native striping capabilities.
storage-tool compose	The <code>raid-c</code> value of the optional argument [-g --geometry]	Specifies the creation of a local RAID C device.	Use array-native striping capabilities.
virtual-volume expand	The option [-e --extent]	The target local device or extent to add to the virtual volume using	Use the <code>storage-volume</code> method of expansion.

Table 2 Deprecated options to commands (continued)

Command	Deprecated options/ arguments	Description	Workaround
		the concatenation method of expansion.	

In addition to these options, the `prod script` that enables slicing-at-the-top functionality will be deprecated in the future VPLEX releases. Use the array-native slicing capabilities to mitigate this deprecation.

CHAPTER 1

Using the VPLEX CLI

- [Logging in to the CLI](#) 22
- [Password Policies](#) 23
- [Logging out from the CLI](#) 24
- [CLI context tree](#) 24
- [Context tree searching](#) 25
- [Navigate the CLI context tree](#) 26
- [Using CLI commands](#) 28

Logging in to the CLI

The VPLEX command line interface supports all VPLEX operations. You can access the CLI in both the VPLEX clusters. Use an SSH client, such as PuTTY and OpenSSH, to log in to the interface.

Before you begin

Before logging in to the CLI, configure the SSH client as follows:

- Port: 22
- SSH protocol version: 2
- Scrollback lines: 20000

Procedure

1. Using an SSH client, connect to the VPLEX management module in the cluster that you want to log in to. Use the following information to connect to the module:
 - The public IP address of MMCS-A, if you are using the VS6 hardware.
 - The public IP address of the management server, if you are using the VS2 hardware.

The following prompt appears:

```
Login as:
```

2. Type `service` to log in as the service user and press ENTER.

The following prompt appears:

```
Using keyboard-interactive authentication.
Password:
```

3. Type the service password and press ENTER.


The following prompt appears:

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

4. Type `vp1excli` and press ENTER.

The following prompt appears:

```
vp1excli:
```

 **Note:** Starting the VPLEX CLI no longer requires a username and password. Please verify that no automated scripts supply usernames or passwords.

Results

You are now logged in to the VPLEX CLI.

Password Policies

The management server uses a Pluggable Authentication Module (PAM) infrastructure to enforce minimum password quality.

For more information about technology used for password protection, refer to the *Security Configuration Guide*.

Note the following:

- Password policies do not apply to users configured using the LDAP server.
- The Password inactive days policy does not apply to the admin account to protect the admin user from account lockouts.
- During the management server software upgrade, an existing user's password is not changed. Only the user's password age information changes.
- You must be an admin user to configure a password policy.

The following table lists and describes the password policies and the default values.

Table 3 Default password policies

Policy name	Description	Default value
Minimum password length	The minimum number of characters used when creating or changing a password. The minimum number of characters includes numbers, upper-case and lower-case letters, and special characters.	8
Minimum password age	The minimum number of days a password can be changed after the last password change.	1 (0 for service account)
Maximum password age	The maximum number of days that a password can be used since the last password change. After the maximum number of days, the account is locked and the user must contact the admin user to reset the password.	90 (3650 days for service account)
Password expiry warning	The number of days before the password expires. A warning message indicating that the password must be changed is displayed.	15 (30 days for service password)
Password inactive days	The number of days after a password has expired before the account is locked.	1

The password policy for existing admin, service, and customer-created user accounts is updated automatically as part of the upgrade to this release. See the *VPLEX Security Configuration Guide* for information about account passwords.

Valid Password Characters

The following characters are allowed in a VPLEXcli password:

- A-Z
- a - z
- 0 - 9
- . ? / * @ ^ % # + = - _ ~ : space

Note: A space is allowed only between the characters in a password, not in the beginning or the end of the password.

Note: A password can not begin with a pound sign (#).

Logging out from the CLI

Use the exit command to exit the command line interface from any context.

About this task

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 command line interface effectively.

The root context contains these sub-contexts:

- `cluster-witness/` - Manage VPLEX Witness options. If the VPLEX Witness optional component is installed, the cluster-witness context will be available.
- `clusters/` - Create and manage links between clusters, devices, extents, system volumes and virtual volumes. Register initiator ports, export target ports, and storage views.
- `connectivity/` - Configure connectivity between back-end storage arrays, front-end hosts, local directors, port-groups and inter-cluster WANs.
- `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 vDirectors, fans, management modules, and power..
- `management-server/` - Manage the Ethernet ports.

- `monitoring/` - Create and manage performance monitors.
- `notifications/` - Create and manage call-home events.
- `recoverpoint/` - Manage RecoverPoint options.
- `security/` - Configure and view authentication password-policy settings. Create, delete, import and export security certificates. Set and remove login banners. The authentication sub context was added to the security context.
- `system-defaults/` - Display systems default settings.

Except for `system-defaults/`, each of the 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 `/`.

Context tree searching

Search the context tree for context names and data matching specific patterns.

Using the Find command to search the context tree

Use this command to find all contexts matching a pattern. When invoked interactively, the command prints the contexts to the screen.

Patterns can be either literal character strings or strings that include wildcard characters. For a complete list of supported CLI wildcard characters, see the topic "Wildcards" in the *VPLEX CLI Reference Guide*.

Example 1 Examples of using the Find command

For example, specifying the literal pattern `-c` returns all contexts containing data matching that specific literal when used with the `*` wildcard. In this case, the command retrieves contexts containing the `-c` option.

```
Vplexcli:/> find -c *
[/alerts, /clusters, /data-migrations, /distributed-storage, /
engines,
/management-server, /monitoring, /notifications, /recoverpoint, /
security,
/system-defaults]
```

Note: `find` searches the contents of the current directory. In the example above, since the current directory is the root, `find` searches the entire context tree.

Using a wildcard returns only results matching a particular string pattern. For example, specifying the pattern `find /clusters/cluster-1/devices/rC*` returns the following contexts matching this pattern.

```
Vplexcli:/> find /clusters/cluster-1/devices/rC*
[/clusters/cluster-1/devices/rC_C1_0000, /clusters/cluster-1/
devices/rC_C1_0001,
```

Example 1 Examples of using the Find command (continued)

```
/clusters/cluster-1/devices/rC_C1_0002, /clusters/cluster-1/devices/
rC_C1_0003]
```

Navigate the CLI context tree

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

The current context is always displayed at the command line interface prompt:

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

For example, to navigate from the root (`/`) context to the `connectivity` context to view member ports for a specified FC port group:

```
VPlexcli:/clusters/cluster-1/connectivity/back-end/port-groups/fc-port-
group-1-0/
member-ports> ll
```

Director	Port	Enabled	Address
director-1-1-A	A1-FC00	enabled	0xc001448783010800
director-1-1-B	B1-FC00	enabled	0xc001448783018800
director-1-2-A	A1-FC00	enabled	0xc001448782da0800
director-1-2-B	B1-FC00	enabled	0xc001448782da8800

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

```
VPlexcli:/> cd clusters/cluster-1/connectivity/back-end/port-groups/iscsi-
port-group-8/member-ports> ll
```

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

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

Use the `cd` command followed by a space and 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-GE01, /]
```

- 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>
```

Where am I in the context tree?

The CLI includes several features to help locate your current position in the context tree and determine what contexts and/or commands are accessible.

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

- The command prompt displays the current context:

```
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
```

- The `cd` command followed by a <Tab> displays the same information as `ls` at the context level.

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

Note: For contexts where the next level down the tree is a list of objects, the `tree` command displays the list. This output can be very long. For example:

Using CLI 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:

- [vplex_c_complete_a_command](#)
- [vplex_c_display_valid_contexts_and_commands](#)

- [vplex_c_display_command_arguments](#)

Complete a command

Use the Tab key to automatically complete a path or command until the path or 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:

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/
```

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

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

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 or contexts for the current context:

About this task

```
Vplexcli:/> cd /clusters/cluster-1/ <Tab>
connectivity/          consistency-groups/
devices/              exports/
performance-policies/ storage-elements/
system-volumes/      uninterruptible-power-supplies/
virtual-volumes/

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 command line interface 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.

 **Note:** Use the `find` command with wildcards to find context names and data matching specific patterns in the CLI context tree. See [Context Tree Searching](#) for more information.

* 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 are named as follows:

Clusters

VPLEX Local configurations have a single cluster, with a cluster ID of cluster 1. VPLEX Metro configurations have two clusters with cluster IDs of 1 and 2.

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

Engines

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

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, those 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

Specifying addresses

VPLEX uses both IPv4 and IPv6 addressing. Many commands can be specified as IPv4 or IPv6 formats.

See the Dell EMC VPLEX Administration Guide for usage rules and address formats.

Command globbing

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

Example 1

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
```

Example 2

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.

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 argument

The `--verbose` argument 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 you type the first letter, 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
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 in the current context, including global commands.
- Use the `help or ?` command with `-G` argument to display all the commands available in 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` to display help for the specified command.

CHAPTER 2

Commands

• advadm dismantle.....	44
• alias.....	45
• amp register.....	47
• amp unregister.....	49
• array claim.....	50
• array forget.....	51
• array re-discover.....	52
• array used-by.....	54
• authentication directory-service configure.....	55
• authentication directory-service map.....	59
• authentication directory-service show.....	60
• authentication directory-service unconfigure.....	62
• authentication directory-service unmap.....	62
• back-end degraded list.....	64
• back-end degraded recover.....	65
• batch-migrate cancel.....	66
• batch-migrate check-plan.....	67
• batch-migrate clean.....	69
• batch-migrate commit.....	70
• batch-migrate create-plan.....	71
• batch-migrate pause.....	73
• batch-migrate remove.....	74
• batch-migrate resume.....	75
• batch-migrate start.....	76
• batch-migrate summary.....	78
• battery-conditioning disable.....	81
• battery-conditioning enable.....	83
• battery-conditioning manual-cycle cancel-request.....	85
• battery-conditioning manual-cycle request.....	86
• battery-conditioning set-schedule.....	87
• battery-conditioning summary.....	90
• cache-invalidate.....	92
• cache-invalidate-status.....	95
• capture begin.....	98
• capture end.....	99
• capture pause.....	100
• capture replay.....	100
• capture resume.....	101
• cd.....	102
• chart create.....	103
• cluster add.....	105
• cluster cacheflush.....	106
• cluster configdump.....	107
• cluster expel.....	109
• cluster forget.....	110

• cluster restart-local-cluster	111
• cluster show-remote-devices	112
• cluster shutdown	114
• cluster status	116
• cluster stop-local-cluster	119
• cluster summary	119
• cluster unexpel	123
• cluster-witness configure	124
• cluster-witness disable	126
• cluster-witness enable	128
• collect-diagnostics	133
• configuration complete-system-setup	135
• configuration configure-auth-service	136
• configuration connect-local-directors	137
• configuration connect-remote-directors	138
• configuration continue-system-setup	139
• configuration cw-vpn-configure	139
• configuration cw-change-password	141
• configuration cw-vpn-reset	141
• configuration enable-front-end-ports	143
• configuration event-notices-reports config	144
• configuration event-notices-reports reset	144
• configuration event-notices-reports-show	145
• configuration flashdir-backup disable	146
• configuration flashdir-backup enable	147
• configuration get-product-type	147
• configuration join-clusters	148
• configuration metadata-backup	149
• configuration register-product	151
• configuration remote-clusters add-addresses	152
• configuration remote-clusters clear-addresses	154
• configuration show-meta-volume-candidates	155
• configuration subnet clear	156
• configuration subnet remote-subnet add	158
• configuration subnet remote-subnet remove	159
• configuration sync-time	161
• configuration sync-time-clear	162
• configuration sync-time-show	163
• configuration system-reset	164
• configuration system-setup	165
• configuration upgrade-meta-slot-count	166
• connect	167
• connectivity director	169
• connectivity list all	170
• connectivity list directors	171
• connectivity list initiators	172
• connectivity list storage-volumes	173
• connectivity show	174
• connectivity validate-be	174
• connectivity validate-local-com	177
• connectivity validate-wan-com	178
• consistency-group add-virtual-volumes	179
• consistency-group choose-winner	180
• consistency-group convert-to-local	183
• consistency-group create	184
• consistency-group destroy	186

• consistency-group list-eligible-virtual-volumes.....	187
• consistency-group remove-virtual-volumes.....	188
• consistency-group resolve-conflicting-detach.....	190
• consistency-group resume-at-loser.....	192
• consistency-group set-detach-rule no-automatic-winner.....	194
• consistency-group set-detach-rule winner.....	195
• consistency-group summary.....	197
• date.....	197
• describe.....	198
• device attach-mirror.....	199
• device collapse.....	201
• device detach-mirror.....	202
• device mirror-isolation auto-unisolation disable.....	206
• device mirror-isolation auto-unisolation enable.....	207
• device mirror-isolation disable.....	209
• device mirror-isolation enable.....	211
• device mirror-isolation show.....	213
• device resume-link-down.....	215
• device resume-link-up.....	217
• device resurrect-dead-storage-volumes.....	218
• director appcon.....	219
• director appdump.....	220
• director appstatus.....	222
• director commission.....	223
• director decommission.....	224
• director fc-port-stats.....	224
• director firmware show-banks.....	226
• director forget.....	227
• director passwd.....	228
• director ping.....	229
• director shutdown.....	230
• director tracepath.....	231
• director uptime.....	232
• dirs.....	233
• disconnect.....	233
• dm migration cancel.....	234
• dm migration clean.....	235
• dm migration commit.....	236
• dm migration pause.....	237
• dm migration remove.....	238
• dm migration resume.....	239
• dm migration start.....	240
• drill-down.....	242
• ds dd convert-to-local.....	244
• ds dd create.....	245
• ds dd declare-winner.....	247
• ds dd destroy.....	248
• ds dd remove-all-rules.....	249
• ds dd set-log.....	250
• ds rule destroy.....	252
• ds rule island-containing.....	252
• ds rule-set copy.....	254
• ds rule-set create.....	255
• ds rule-set destroy.....	255
• ds rule-set what-if.....	257
• ds summary.....	258

- [esrs import-certificate](#) 265
- [esrs register](#) 266
- [esrs status](#) 267
- [esrs un-register](#) 268
- [event-test](#) 269
- [exec](#) 270
- [exit](#) 270
- [export initiator-port discovery](#) 271
- [export initiator-port register](#) 272
- [export initiator-port register-host](#) 273
- [export initiator-port show-logins](#) 274
- [export initiator-port unregister](#) 276
- [export port summary](#) 277
- [export storage-view addinitiatorport](#) 279
- [export storage-view addport](#) 280
- [export storage-view addvirtualvolume](#) 280
- [export storage-view checkconfig](#) 283
- [export storage-view create](#) 284
- [export storage-view destroy](#) 285
- [export storage-view find](#) 286
- [export storage-view find-unmapped-volumes](#) 287
- [export storage-view map](#) 288
- [export storage-view removeinitiatorport](#) 289
- [export storage-view removeport](#) 290
- [export storage-view removevirtualvolume](#) 291
- [export storage-view show-powerpath-interfaces](#) 292
- [export storage-view summary](#) 292
- [export target-port renamewwns](#) 293
- [extent create](#) 294
- [extent destroy](#) 296
- [extent summary](#) 297
- [find](#) 299
- [front-end-performance-stats start](#) 300
- [front-end-performance-stats status](#) 300
- [front-end-performance-stats stop](#) 301
- [getsysinfo](#) 302
- [health-check](#) 303
- [help](#) 309
- [history](#) 310
- [iscsi chap back-end add-credentials](#) 311
- [iscsi chap back-end disable](#) 312
- [iscsi chap back-end enable](#) 313
- [iscsi chap back-end list-credentials](#) 314
- [iscsi chap back-end remove-credentials](#) 315
- [iscsi chap back-end remove-default-credential](#) 316
- [iscsi chap back-end set-default-credential](#) 317
- [iscsi chap front-end add-credentials](#) 318
- [iscsi chap front-end disable](#) 320
- [iscsi chap front-end enable](#) 321
- [iscsi chap front-end list-credentials](#) 322
- [iscsi chap front-end remove-credentials](#) 323
- [iscsi chap front-end remove-default-credential](#) 324
- [iscsi chap front-end set-default-credential](#) 325
- [iscsi check-febe-connectivity](#) 326
- [iscsi isns add](#) 327
- [iscsi isns list](#) 328

• iscsi isns remove	330
• iscsi sendtargets add	331
• iscsi sendtargets list	332
• iscsi sendtargets rediscover	334
• iscsi sendtargets remove	335
• iscsi targets list	336
• iscsi targets logout	337
• license install	338
• license reset	339
• license show	340
• local-device create	341
• local-device destroy	344
• local-device summary	345
• log filter create	346
• log filter destroy	348
• log filter list	348
• log source create	349
• log source destroy	350
• log source list	351
• logging-volume add-mirror	351
• logging-volume create	352
• logging-volume detach-mirror	357
• logging-volume destroy	359
• logical-unit forget	360
• ls	361
• management-server set-ip	364
• manifest upgrade	366
• manifest version	366
• meta-volume attach-mirror	367
• meta-volume backup	368
• meta-volume create	369
• meta-volume destroy	372
• meta-volume detach-mirror	373
• meta-volume move	374
• meta-volume verify-on-disk-consistency	375
• monitor add-console-sink	376
• monitor add-file-sink	378
• monitor collect	379
• monitor create	380
• monitor destroy	383
• monitor get-stats	384
• monitor remove-sink	385
• monitor stat-list	385
• ndu pre-check	386
• ndu pre-config-upgrade	387
• ndu recover	388
• ndu rolling-upgrade ssd-fw	388
• ndu start	391
• ndu status	394
• notifications call-home import-event-modifications	395
• notifications call-home remove-event-modifications	395
• notifications call-home view-event-modifications	396
• notifications call-home test	397
• notifications job cancel	398
• notifications job delete	398
• notifications job resubmit	400

- [notifications snmp-trap create](#)..... 401
- [notifications snmp-trap destroy](#)..... 402
- [password-policy reset](#)..... 402
- [password-policy set](#)..... 403
- [plugin addurl](#)..... 405
- [plugin listurl](#)..... 406
- [plugin register](#)..... 407
- [popd](#)..... 407
- [pushd](#)..... 408
- [rebuild set-transfer-size](#)..... 409
- [rebuild show-transfer-size](#)..... 410
- [rebuild status](#)..... 411
- [report aggregate-monitors](#)..... 412
- [report capacity-arrays](#)..... 413
- [report capacity-clusters](#)..... 416
- [report capacity-hosts](#)..... 417
- [report create-monitors](#)..... 419
- [report poll-monitors](#)..... 421
- [rm](#)..... 422
- [rp import-certificate](#)..... 423
- [rp rpa-cluster add](#)..... 424
- [rp rpa-cluster remove](#)..... 431
- [rp summary](#)..... 432
- [rp validate-configuration](#)..... 434
- [schedule add](#)..... 437
- [schedule list](#)..... 438
- [schedule modify](#)..... 438
- [schedule remove](#)..... 439
- [scheduleSYR add](#)..... 440
- [scheduleSYR list](#)..... 441
- [scheduleSYR remove](#)..... 442
- [script](#)..... 442
- [security configure-mmcs-users](#)..... 443
- [security create-ca-certificate](#)..... 444
- [security create-certificate-subject](#)..... 446
- [security create-host-certificate](#)..... 448
- [security configure-certificates](#)..... 450
- [security delete-ca-certificate](#)..... 454
- [security delete-host-certificate](#)..... 455
- [security export-ca-certificate](#)..... 456
- [security export-host-certificate](#)..... 457
- [security import-ca-certificate](#)..... 458
- [security import-host-certificate](#)..... 459
- [security ipsec-configure](#)..... 461
- [security list-certificates](#)..... 463
- [security remove-login-banner](#)..... 464
- [security renew-all-certificates](#)..... 465
- [security set-login-banner](#)..... 469
- [security show-cert-subj](#)..... 470
- [security web-configure](#)..... 471
- [sessions](#)..... 471
- [set](#)..... 472
- [set topology](#)..... 477
- [show-use-hierarchy](#)..... 478
- [sms dump](#)..... 479
- [snmp-agent configure](#)..... 481

• snmp-agent start	482
• snmp-agent status	483
• snmp-agent stop	483
• snmp-agent unconfigure	484
• source	484
• storage-tool dismantle	485
• storage-tool compose	486
• storage-volume auto-unbanish-interval	489
• storage-volume claim	490
• storage-volume claimingwizard	493
• storage-volume find-array	496
• storage-volume forget	498
• storage-volume list-banished	499
• storage-volume list-thin-capable	500
• storage-volume resurrect	501
• storage-volume summary	503
• storage-volume unbanish	507
• storage-volume unclaim	508
• storage-volume used-by	510
• syrcollect	511
• tree	512
• unalias	513
• user add	514
• user event-server add-user	515
• user event-server change-password	516
• user list	517
• user passwd	518
• user remove	519
• user reset	520
• validate-system-configuration	521
• vault go	522
• vault overrideUnvaultQuorum	523
• vault status	525
• verify fibre-channel-switches	529
• version	529
• virtual-volume create	533
• virtual-volume destroy	540
• virtual-volume expand	541
• virtual-volume list-thin	545
• virtual-volume provision	546
• virtual-volume re-initialize	548
• virtual-volume set-thin-enabled	548
• virtual-volume summary	549
• vpn restart	553
• vpn start	553
• vpn status	554
• vpn stop	555
• wait	555
• webserver	556

advadm dismantle

Dismantles 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	
<code>[-r --devices] <i>context path,context path...</i></code>	One or more devices to dismantle. Entries must be separated by commas. You can use glob patterns.
<code>[-v --virtual-volumes] <i>context path,context path...</i></code>	One or more virtual volumes to dismantle. Entries must be separated by commas. You can use glob patterns.
Optional Arguments	
<code>--unclaim-storage-volumes</code>	Unclaim the storage volumes after the dismantle is completed.
<code>[-f --force]</code>	Force the dismantle without asking for confirmation. Allows the command to be run from a non-interactive script.

Description

To dismantle a virtual volume, the specified volume must:

- Not be exported to a storage view.
- Not a member of a consistency group

virtual volume exported through a storage view or belonging to a consistency group are not eligible to be dismantled. The command skips any volumes that are not eligible for dismantle, prints a message listing skipped volumes, and dismantles those volumes that are eligible.

If the `--force` argument is used, no confirmation is displayed before the dismantle.

Examples

In the following example, the specified volume is dismantled:

```
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
```

In the following example, the specified volumes are NOT dismantled because they are exported or are members of a consistency group:

```

Vplexcli:/>advadm dismantle -v rC_extentSrc_C1_CHM_00*, axel_dr1_vol
The following virtual-volumes will not be dismantled because they are
exported. Please remove
them from the storage-views before dismantling them:
    /clusters/cluster-1/virtual-volumes/rC_extentSrc_C1_CHM_0002_vol is in
/clusters/cluster-1/exports/storage-views/chimera_setupTearDown_C1
    /clusters/cluster-1/virtual-volumes/rC_extentSrc_C1_CHM_0001_vol is in
/clusters/cluster-1/exports/storage-views/chimera_setupTearDown_C1
.
.
.
The following virtual-volumes will not be dismantled because they are in
consistency-groups.
Please remove them from the consistency-groups before dismantling them:
    /clusters/cluster-2/virtual-volumes/axel_dr1_vol is in
/clusters/cluster-2/consistency-groups/async_sC12_vC12_nAW_CHM

No virtual-volumes to dismantle.

```

See also

- ds dd create
- iscsi sendtargets add
- virtual-volume create

alias

Creates a command alias.

Contexts

All contexts.

Syntax

alias

[-n|--name] *name*

[-t|--to] "*commands and arguments*"

Arguments

Required arguments	
[-n --name] <i>name</i>	<p>* The name of the new alias.</p> <ul style="list-style-type: none"> • Up to 63 characters. • May contain letters, numbers, and underscores '_'. s • Cannot start with a number.
[-t --to] " <i>commands and arguments</i> "	<p>* A string of commands and arguments enclosed in quotation marks. This string is invoked when the aliased command is used.</p>

* - 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 system.

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.

- `? 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).

The following aliases are pre-configured:

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

An `alias` set at the command line does not persist when the user interface is restarted. To create an `alias` command that persists, add it to the `/var/log/VPlex/cli/VPlexcli-init` file.

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

Examples

Create an alias:

```
VPlexcli:/> 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`
- `unalias`

amp register

Associates an Array Management Provider to a single Cluster.

Contexts

`/clusters/ClusterName/storage-elements/array-providers`

Syntax

```
amp register
[-n | --name] name
[-i | --ip-address] ip-address of Array Management Provider
[-a | --array] =storage-array
[-c | --cluster] cluster context path
[-s | --use-ssl]
[-h | --help]
[--verbose]
[-u | --username] user name
[-t | --provider-type] provider type
[-p | --port-number] port number
```

Arguments

Required arguments	
<code>[-n --name] <i>name</i></code>	* The name of the Array Management Provider.
<code>[-i --ip-address] <i>ip-address of Array Management Provider-</i></code>	* The IP address of the Array Management Provider.
<code>[-u --username] <i>user name</i></code>	* The user name required for connecting to the Array Management Provider.
<code>[-t --provider-type] <i>provider type</i></code>	* The type of Array Management Provider. SMI-S and REST are currently the only provider types supported.
<code>[-p --port-number] <i>port number</i></code>	* The port number to use along with the IP-address to construct the URL indicating the Array Management Provider.
Optional arguments	

<code>[-a --array]=<i>storage-array</i></code>	If the provider-type is REST, this argument specifies the name of the array being registered. Note that this option is deprecated, using it will have no effect.
<code>[-c --cluster] <i>cluster context path</i></code>	The VPLEX cluster associated with the Array Management Provider. This argument may be omitted when the command is executed from or below a cluster context (meaning the cluster is implied).
<code>[-s --use-ssl]</code>	Specifies whether to use HTTP or HTTPS (SSL) protocol when connecting to the Array Management Provider URL.
<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution.

* - positional argument

Description

An Array Management Provider (AMP) is an external array management endpoint that VPLEX communicates with to execute operations on individual arrays. Examples of AMPs include external SMI-S and REST providers.

An AMP exposes a management protocol/API. Array operations can be executed through this API. Examples of Management protocols include Provisioning and Snap & Clone.

An AMP manages one or more arrays. For example, an SMI-S provider can manage multiple arrays.

Examples

Registering (adding) an array provider:

```
Vplexcli:/> amp register -n mysmis -i 10.12.234.12 -u admin -t SMI-S -p 1234 -c Hopkinton
Enter array provider password:
Enter array provider password again for verification:
Successfully registered array provider.
```

Registering (adding) an array provider with an invalid provider type:

```
Vplexcli:/> amp register -n mysmis -i 10.12.234.12 -u admin -t vasa -p 1234 -c Hopkinton
amp register: Evaluation of <<amp register -n mysmis -i 10.12.234.12 -u admin -t vasa -p 1234 -c Hopkinton>> failed.
cause: Command execution failed.
cause: Failed to set value for option --provider-type.
cause: 'vasa' is not a valid amp-type for command 'amp register'.
Valid values: 'smi-s' (not case sensitive).
```

See also

- amp unregister

amp unregister

Unregisters Array Management Provider. The Array Management provider is no longer available for any operations after being unregistered.

Contexts

/clusters/ClusterName/storage-elements/array-providers

Syntax

```
amp unregister
[-n | --name] name
[-c | --cluster] cluster context
[-f | --force]
[-h | --help]
[--verbose]
```

Arguments

Required arguments	
<code>[-n --name] <i>name</i></code>	* The name of the array management provider to unregister.
Optional arguments	
<code>[-c --cluster] <i>cluster context</i></code>	The cluster associated with the Array Management Provider. Can be omitted when the command is executed from or below a cluster context in which case the cluster is implied.
<code>[-f --force]</code>	Force the operation without confirmation. Allows the command to be run from a non-interactive script.
<code>[-h --help]</code>	Displays the usage for this command.
<code> [--verbose]</code>	Provides more output during command execution.

* - positional argument

Description

An Array Management Provider (AMP) is an external array management endpoint that VPLEX communicates with to execute operations on individual arrays. Examples of AMPs include external SMI-S providers

The `amp unregister` command unregisters an array provider. The AMP is no longer available for any operations.

After being unregistered, the array-provider is no longer available. Re-registering the provider results in definition of a new provider.

Examples

Unregistering an array provider:

```
Vplexcli:/> amp unregister -n myAmpThatExistsAndIsNotFake
Warning: Unregistering an array-provider will permanently remove it.
```

```
Do you wish to proceed? (Yes/No): Yes
VPlexcli:/>
```

Unregistering an array provider with an invalid provider:

```
VPlexcli:/> amp unregister -n myExcellentAMP -c cluster-1
Unregistering an array-provider will permanently remove it. Do you wish to
proceed? (Yes/No) Yes
amp unregister: Evaluation of <<amp unregister -n myExcellentAMP -c
cluster-1>> failed.
cause: Failed to unregister array-provider 'myExcellentAMP' on
cluster 'cluster-1'.
cause: The array-provider 'myExcellentAMP' is not a registered
array-provider on cluster 'cluster-1'.
VPlexcli:/>
```

See also

- amp register

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]
[--force]
```

Arguments

Required arguments	
<code>[-s --storage-array] <i>context-path</i></code>	* Context path of the storage-array on which to claim storage volumes.
Optional arguments	
<code>[-m --mapping-file] <i>mapping file</i></code>	Location of the name mapping file.
<code>[-t --tier] <i>mapping file</i></code>	Add a tier identifier to the storage volumes to be claimed.
<code>[-l --claim]</code>	Try to claim unclaimed storage-volumes.
<code> [--force]</code>	Force the operation without confirmation. Allows the command to be run from a non-interactive script.

* - argument is positional.

Description

Claims and names unclaimed storage volumes for a given array.

Some storage arrays support auto-naming (Dell EMC Symmetrix/VMAX, CLARiiON/VNX, XtremIO, Hitachi AMS 1000, 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.

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.

This command can fail if there is not a sufficient number of meta-volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

See also

- `storage-volume find-array`

array forget

Removes a storage-array that is being retired from VPLEX.

Context

All contexts.

Syntax

```
array forget [-h|--help]
             [--verbose]
             [-r|--retire-logical-units]
             [-a|--array]array
```

Arguments

Optional arguments	
<code>-h --help</code>	Displays the usage for this command.
<code>--verbose</code>	Provides more output during command execution. This might not have any effect for some commands.
<code>-r --retire-logical-units</code>	Retires all logical units before retiring the array. If not specified, the command fails if there are still logical units from the array in the logical-units context on VPLEX.
Required arguments	
<code>-a --array= <i>array</i></code>	Specifies the context path of the storage-array to forget.

* - argument is positional

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
[-d|--hard]
[-f|--force]
```

Arguments

Required arguments	
<code>[-a --array] <i>context-path</i></code>	* Context path that specifies the storage-array to re-discover.
<code>[-c --cluster] <i>cluster-id</i></code>	Cluster ID of the target cluster.
Optional arguments	
<code>[-d --hard]</code>	<ul style="list-style-type: none"> Perform a hard rediscover. This is a disruptive operation because ITLs are destroyed and full discoveries executed. I/O temporarily stops until the array responds with data for each LUN. Discovery time correlates to array response time, number of provisioned volumes, and number of paths per volume. Large numbers of volumes result in longer discovery times. VPLEX automatically verifies the volume ID (VPD ID) on existing provisioned volumes to detect if the array's device/LUN mapping has changed. *LUN swapping: Logical-unit swapping occurs when the array's back-end device/LUN mapping has changed. This can be detected by comparing the system's saved copy of the volume's ID (VPD_ID) with value returned by INQ VPD83 to its LUN. For example: A LUN is removed from a storage group on an array and then re-added. The LUN may now be mapped to a different device which reports a different VPD_ID value. Data corruption could occur if writes are sent to old VPD_ID value. If logical-unit swapping has occurred use the <code>--hard</code> option to force fresh discovery of all ITLs on the array. <p> Note: using the <code>--hard</code> option is disruptive and can result in data unavailability and/or data loss on live exported paths.</p>
<code>[-f --force]</code>	Force the operation without confirmation. Allows the command to be run from a non-interactive script.

from the individual storage array context:

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

See also

- storage-volume find-array

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] <i>context-path</i>	* 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.

Examples

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_6006016061211100d42febbalbade011_1:
  extent_6006016061211100d42febbalbade011_1
  VPD83T3:6006016061211100d42febbalbade011
/distributed-storage/distributed-devices/dev1_source:
  dev1_source2012Feb16_191413
  extent_sv1_1
```

```

svl
/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

```

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`
- `storage-volume summary`

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"
[-n|--bind-distinguished-name] "bind-distinguished-name"
[-l] "pathname of certificate file"
[-p|--bind-password]
[-r|--user-search-path] "user-search-path"
[-s|--server-name] server-name
[-t|--connection-type] [1|2]
[-o|--port] port
[-c|--custom-attributes]
[-u|--map-user] "map-user"
[-g|--map-group] "map-group"
--dry-run

```

Arguments

Note: The `-m` option is no longer supported. Use the `-g` and `-u` options for managing access to the management server by groups and users.

Required arguments	
<code>[-d --directory-server] [1 2]</code>	<p>Specifies the directory server to configure on the cluster to authenticate users.</p> <ul style="list-style-type: none"> 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. <p>Note: If option 2 (Active Directory) is selected, use the <code>--custom-attributes</code> argument to map attributes if directory server UNIX attributes are different from the default attributes mapped by VPLEX.</p>
<code>[-i --ip-address] <i>IP address</i></code>	IP address of the directory server.
<code>[-b --base-distinguished-name] "<i>base-distinguished-name</i>"</code>	<p>The base Distinguished Name (DN) of the directory server. A DN is a sequence of relative distinguished names connected by commas. The Distinguished Name must be enclosed in quotes. For example:</p> <pre>dc=org,dc=company,dc=com</pre>
<code>[-n --bind-distinguished-name] "<i>bind-distinguished-name</i>"</code>	<p>The Bind Distinguished Name of the OpenLDAP/Active Directory server. The Bind Distinguished Name must be enclosed in quotes. For example:</p> <pre>cn=Manager,dc=my-domain,dc=com</pre>
<code>[-l] "<i>pathname of certificate file</i>"</code>	Configure a certificate.
<code>[-p] --bind-password]</code>	Password of Bind Distinguished Name. A prompt for the password is displayed. The password is not displayed as it is typed.
<code>[-r --user-search-path] "<i>user-search-path</i>"</code>	The distinguished name of the node at which to begin user searches in the directory server. <i>user-search-path</i> is a sequence of relative distinguished names connected by commas. The principal specified for a <i>user-search-path</i> must belong to an organizational unit class. The <i>user-search-path</i> must be enclosed in quotes.
Optional arguments	
<code>[-s --server-name] <i>server-name</i></code>	Name of the directory server. This argument is required when the <code>--connection-type</code> argument is specified with a value of 2 (LDAPS).
<code>[-t --connection-type] {1 2}</code>	Select the cryptographic protocol to use to connect to the LDAP/Active Directory server. Values are:

	<ul style="list-style-type: none"> • 1 - (Default) - Use LDAP. Communication with the LDAP/Active Directory server will be in plain text. • 2 - Use the Secure LDAP (LDAPS). If LDAPS is selected, use the <code>--server-name</code> argument to specify the LDAP server.
<code>[-o --port] <i>port</i></code>	Port number of the LDAP/Active Directory server. Range: 1-65536. Default: 389 when <code>--connection-type</code> is set to LDAP (1).
<code>[-c --custom-attributes]</code>	Provide custom attribute names for attribute mapping. Prompts for mapping the attribute names.
<code>[-u --map-user] "<i>map-user</i>"</code>	Specifies which users can login to the management server. The <code>map-user</code> must be enclosed in quotes.
<code>[-g --map-group] "<i>map-group</i>"</code>	Specifies a group. Only members of this group within the user search path can log in to the management server. The <code>map-group</code> must be enclosed in quotes.
<code>--dry-run</code>	Run the command but don't do anything.

Description

This command configures an authentication service on the VPLEX cluster.

VPLEX supports two types of 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.

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.

In order to authenticate directory service users, 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 LDAP or LDAPS 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
```

Examples

Configure the Active Directory directory service on the VPlex cluster:

```
VPlexcli:/> authentication directory-service configure -d 2 -i 192.168.98.101
-b "dc=org,dc=company,dc=com" -r
"ou=vplex,dc=org,dc=company,dc=com" -n
"cn=Administrator,cn=Users,dc=org,dc=company,dc=com" -t 2 -s servername -p
```

OR

```
VPlexcli:/> authentication directory-service configure -d 1 -i 192.168.98.102
-b "dc=org,dc=company,dc=com" -r
"ou=vplex,dc=org,dc=company,dc=com" -g
"cn=testgroup,dc=org,dc=company,dc=com" -n
"cn=Administrator,dc=org,dc=company,dc=com" -o
9389 -p
```

OR

```
VPlexcli:/> authentication directory-service configure -d 2 -i 192.168.98.101
-b "dc=org,dc=company,dc=com" -r
"ou=vplex,dc=org,dc=company,dc=com" -u
"uid=testuser,ou=vplex,dc=org,dc=company,dc=com" -n
"cn=Administrator,cn=Users,dc=org,dc=company,dc=com" --custom-attributes -p
```

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**. Values for custom attributes are case sensitive. Verify the case when specifying the value for a custom attribute.

```
VPlexcli:/> authentication directory-service configure -d 2 -i 10.31.52.189 -
b "dc=intranet,dc=vplex,dc=com" -r
"ou=qe,ou=vbu,ou=glb,dc=intranet,dc=vplex,dc=com" -u
"cn=user1,ou=qe,ou=vbu,ou=glb,dc=intranet,dc=vplex,dc=com" -g
"cn=devgroup,ou=qe,ou=vbu,ou=glb,dc=intranet,dc=vplex,dc=com" -n
"cn=Administrator,cn=Users,dc=intranet,dc=vplex,dc=com" -p --custom-
attributes
Please note that configuring LDAP means that network channel data is
unencrypted. A better option is to use LDAPS so the data is encrypted. Are
you sure you want to continue the configuration with LDAP. Continue? (Yes/No)
Yes
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]: primaryGroupID
Set value for loginShell attribute [msSFU30LoginShell]: LoginShell
Set value for homeDirectory attribute [msSFU30HomeDirectory]:
unixHomeDirectory
Connecting to authentication server (may take 3 minutes) ...
VPlexcli:/>
```

Configure the OpenLDAP service on the management server:

Note: To define a different posixGroup attribute use custom attributes. Use an appropriate objectClass attribute for posixGroup (e.g., “posixGroup” or “groupOfNames” or “Group”) as used by the OpenLDAP server.

```
Vplexcli:/> authentication directory-service configure
-d 1
-i "3ffe:80c0:22c:803c:211:43ff:fede:9798"
-b "dc=security,dc=sve,dc=emc,dc=com"
-r "ou=qe,dc=security,dc=sve,dc=emc,dc=com"
-n "cn=Administrator,dc=security,dc=sve,dc=emc,dc=com"
-p
Enter Administrator's password:
Connecting to authentication server (may take 3 minutes) ...
Vplexcli:/>
```

See also

- authentication directory-service map
- authentication directory-service show
- authentication directory-service unconfigure

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"
[-u|--map-user] "user-principal"
[-g|--map-group] "group-principal"
--dry-run
```

Arguments

Optional arguments	
<pre>[-m --map-principal] "<i>map-principal</i>"</pre>	<p>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:</p> <pre>OU=eng,dc=vplex,dc=security,dc=lab,dc=emc,dc=com</pre> <p>Note: Option <code>-m</code> for commands can only be used with the older configuration. Use the <code>-g</code> and <code>-u</code> options for adding user and group access to the management server.</p>
<pre>[-u --map-user] "<i>user-principal</i>"</pre>	<ul style="list-style-type: none"> • The map-principal must be enclosed in quotes.

	<ul style="list-style-type: none"> To include a backslash (\) in the map-principal, precede the backslash by a second backslash. For example, to specify the map-principal: <pre>CN=user\ test: (IT),OU=group,DC=company,DC=com</pre> <p>Enter:</p> <pre>"CN=user\\ test: (IT),OU=group,DC=company,DC=com"</pre> <ul style="list-style-type: none"> Users matching this search string will be authenticated. The user-principal must be enclosed in quotes.
<pre>[-g --map-group] "group-principal"</pre>	<p>Users in the group matching this search string are authenticated. Members of the group should be part of the user-search-path specified during the configuration. The <i>group-principal</i> must be enclosed in quotes.</p>
<pre>--dry-run</pre>	<p>Run the command but don't do anything.</p>

Description

A directory server user is an account that can log in to the VPLEX Management Server. Users can be specified explicitly or implicitly via groups.

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.

Examples

Map an LDAP user to a VPLEX cluster:

```
Vplexcli> authentication directory-service map -u
"cn=user2,ou=qe,ou=vbu,ou=glb,dc=intranet,dc=vplex,dc=com"
```

Map an LDAP user group to a VPLEX cluster:

```
Vplexcli> authentication directory-service map -g
"cn=unix_group,dc=intranet,dc=vplex,dc=com"
```

See also

- authentication directory-service configure
- authentication directory-service show
- authentication directory-service unmap

authentication directory-service show

Displays configuration of the directory service used to authenticate users.

Contexts

All contexts.

Syntax

```
authentication directory-service show
```

Description

The fields shown in the output of the `authentication directory-service show` command are described in the following table.

Table 4 authentication directory-service show field descriptions

Field	Description
<code>default-authentication-service</code>	Always Native VPLEX
<code>external-authentication-service</code>	Directory service configured on the cluster. <ul style="list-style-type: none"> OpenLDAP - Lightweight Directory Access Protocol. AD - Windows Active Directory 2003 with SFU 3.5
<code>ip</code>	IP address of the directory server.
<code>base-dn</code>	Base Distinguished Name of the directory service. Components: <ul style="list-style-type: none"> dc - domainComponent cn - commonName ou - organizationalUnitName uid - userid
<code>connection-type</code>	Cryptographic protocol used to connect to the LDAP server. <ul style="list-style-type: none"> LDAP - Plain text protocol LDAPS - Secure protocol
<code>uri</code>	Uniform resource identifier of the directory server.
<code>mapped-principal</code>	Users and/or groups mapped to directory server.

Examples

Display the authentication service configuration:

```
Vplexcli:/> authentication directory-service show
external-authentication-service: OpenLDAP
ip: 10.31.52.53
User Search Path: ou=qe,dc=security,dc=sve,dc=emc,dc=com
connection-type: Plain Text TLS
mapped-principal:
    [Users ]:
    ['uid=testUser3,ou=qe,dc=security,dc=sve,dc=emc,dc=com']
```

See also

- authentication directory-service configure
- authentication directory-service map

authentication directory-service unconfigure

Removes the configuration of a directory service on the VPLEX cluster.

Contexts

All contexts.

Syntax

```
authentication directory-service unconfigure
[-f|--force]
```

Arguments

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

Description

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

Examples

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
- authentication directory-service show
- authentication directory-service unmap

authentication directory-service unmap

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


Contexts

All contexts.

Syntax

```
authentication directory-service unmap
[-m|--map-principal] "mapped-principal"
[-u|--map-user] "user-principal"
[-g|--map-group] "group-principal"
```

Arguments

 **Note:** Option `-m` for commands can only be used with the older configuration. Use the `-g` and `-u` options for adding user and group access to the management server.

Optional arguments	
<pre>[-m --map-principal] "<i>mapped-principal</i>"</pre>	<p>Mapped directory server Distinguished Name to unmap. For example:</p> <pre>ou=eng,dc=vplex,dc=security,dc=lab,dc=emc,dc=com</pre> <ul style="list-style-type: none"> The <i>map-principal</i> 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: <pre>CN=user\ test: (IT),ou=group,DC=company,DC=com</pre> <p>Enter:</p> <pre>"CN=user\\ test: (IT),OU=group,DC=company,DC=com"</pre> <p>Note: The "<i>mapped-principal</i>" argument is case-sensitive and varies depending on the directory service configured:</p> <p>If the directory service is OpenLDAP, entries in the mapped principal argument ('ou', 'uid', 'dc') must be specified in small letters (no capitals).</p> <p>If the directory service is AD, entries in the mapped principal argument ('OU', 'CN', 'DC') must be specified in capital letters.</p> <p>Use the authentication directory-service show command to display the current mapped principals.</p>
<pre>[-u --map-user] "<i>user-principal</i>"</pre>	<p>Users matching this search string will be authenticated. The <i>user-principal</i> must be enclosed in quotes.</p>
<pre>[-g --map-group] "<i>group-principal</i>"</pre>	<p>Users in the group matching this search string are authenticated. Members of the group should be part of the <i>user-search-path</i> specified during the configuration. The <i>group-principal</i> must be enclosed in quotes.</p>

Description

This command unmaps a directory server user or a user group from the VPLEX cluster. There must be at least one principal mapped. If there is only one principal, it cannot be unmapped. If a user search path is specified, unmapping all users and group principals will provide access to all users in the user search path.

CAUTION VPLEX does not currently support unmapping LDAP users from LDAP servers if the users have been removed from the servers. Ensure that the LDAP user is unmapped from VPLEX before removing the user from the LDAP server.

Examples

Unmap a directory server user group from a VPLEX cluster:

```
Vplexcli:/> authentication directory-service unmap -g
"cn=testgroup,dc=org,dc=company,dc=com"
```

OR

```
Vplexcli:/> authentication directory-service unmap -u
"uid=testuser,ou=vplex,dc=org,dc=company,dc=com"
```

OR

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

See also

- authentication directory-service configure
- authentication directory-service map
- authentication directory-service show

back-end degraded list

Displays a list of currently degraded I-Ts.

Contexts

All contexts.

Syntax

back-end degraded list

[-h|--help]

[--verbose]

[-g|--group-by= <group_by>]

Arguments

Optional arguments	
[-h --help]	Display the usage for this command.
[--verbose]	Provides more output during command execution.
[-g --group-by= <group_by>]	Group degraded I-Ts by the specified field. Supported fields: array, director.

Description

Lists I-Ts that have degraded performance, and I-Ts that have been isolated manually or isolated due to unstable performance.

Examples

List all degraded I-Ts grouped by director.

```
VPlexcli:/> back-end degraded
list

Degraded I-Ts:

Director          Director Port Initiator          Target
Array            Degradation Reason
-----
director-1-1-A A1-FC00          0xc00144878bda0800 0x5000144260321e00 EMC-
Invista-rc-surry-1 Manually isolated
                    A1-FC01          0xc00144878bda0900 0x5006016547e01af9 EMC-
CLARiiON-APM00164919257 Manually isolated
director-1-1-B B1-FC01          0xc00144878bda8900 0x5006016547e01af9 EMC-
CLARiiON-APM00164919257 Manually isolated
```

List all degraded I-Ts grouped by array

```
VPlexcli:/> back-end degraded list --group-by
array

Degraded I-Ts:

Array          Director          Director Port Initiator
Target        Degradation Reason
-----
EMC-CLARiiON-APM00164919257 director-1-1-A A1-FC01          0xc00144878bda0900
0x5006016547e01af9 Manually isolated
                    director-1-1-B B1-FC01          0xc00144878bda8900
0x5006016547e01af9 Manually isolated
EMC-Invista-rc-surry-1 director-1-1-A A1-FC00          0xc00144878bda0800
0x5000144260321e00 Manually isolated
```

See also

back-end degraded recover

back-end degraded recover

Recovers the specified degraded I-Ts.

Contexts

All contexts.

Syntax

```
back-end degraded recover
```

```
[-h|--help]
```

```
[--verbose]
```

```
[-p|--paths= <paths>]
```

```
[--all]
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Display the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution.
<code>[-p --paths=<paths>]</code>	The degraded I-Ts to recover. Each I-T must be expressed as a pair in the form " <code>(<initiator>,<target>)</code> ".
<code>[--all]</code>	Recover all currently degraded I-Ts.

Description

Assert that the specified I-Ts are healthy and move them out of their degraded state.

Examples

Recover a specific degraded I-T.

```
VPlexcli:/> back-end degraded recover -p
(0xc00144878bda0900,0x5006016547e01af9)

Recovered I-Ts:

Director      Director Port Initiator      Target
Array                Degradation Reason
-----
director-1-1-A A1-FC01      0xc00144878bda0900 0x5006016547e01af9 EMC-
CLARiiON-APM00164919257 Manually isolated
```

Recover all degraded I-Ts.

```
VPlexcli:/> back-end degraded recover --
all

Recovered I-Ts:

Director      Director Port Initiator      Target
Array                Degradation Reason
-----
director-1-1-A A1-FC00      0xc00144878bda0800 0x5000144260321e00 EMC-
Invista-rc-surry-1 Manually isolated
director-1-1-B B1-FC01      0xc00144878bda8900 0x5006016547e01af9 EMC-
CLARiiON-APM00164919257 Manually isolated
```

See also

`back-end degraded list`

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] <i>pathname</i>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.

Description

Attempts to cancel every migration in the specified batch file. If the command encounters an error, the command prints a warning to the console and continues until every migration listed in the file has been processed.

Note: In order to re-run a canceled migration plan, first run the `batch-migrate remove` command to remove the records of the migration.

Examples

The following shows an example of the `batch-migrate cancel` command used to cancel every migration in the `migrate.txt` file.

```
VPlexcli:/data-migrations/device-migrations>
batch-migrate cancel --file migrate.txt
```

See also

- `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`

batch-migrate check-plan

Checks a batch migration plan.

Contexts

All contexts.

Syntax

```
batch-migrate check-plan
```

```
[-f|--file] pathname
```

Arguments

Required arguments	
<code>[-f --file]</code> <i>pathname</i>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> 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

Check all migration plans before beginning execution.

Examples

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.
```

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`
- `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`

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	
<code>[-f --file] <i>pathname</i></code>	*Directory and filename of migration plan file. relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.
Optional arguments	
<code>[-e --rename-targets]</code>	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`
- `batch-migrate check-plan`
- `batch-migrate commit`
- `batch-migrate create-plan`
- `batch-migrate pause`
- `batch-migrate remove`

batch-migrate commit

Commits the specified batch migration.

Contexts

All contexts.

Syntax

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

Arguments

Required argument	
<code>[-f --file] <i>pathname</i></code>	*Directory and filename of migration plan file. relative paths can be used. if no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> 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 the command encounters an error, the command displays a warning 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.

In order to clean a migration job, you must first commit the job.

Use the `batch-migrate summary` command to verify that the migration has completed with no errors before committing the migration.

Examples

This example commits a list of batch migrations specified in `BSO_19`.

```
VPlexcli:/> batch-migrate commit --file BSO_19
```

See also

- `batch-migrate cancel`
- `batch-migrate check-plan`
- `batch-migrate clean`
- `batch-migrate create-plan`
- `batch-migrate remove`

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	
<code>[-f --sources] <i>local-devices</i></code>	* List of <i>local-devices</i> to migrate virtual volumes from. May contain wildcards.
<code>[-t --targets] <i>local-devices</i></code>	* List of <i>local-devices</i> to migrate the source virtual volumes to. May contain wildcards.
<code>--file <i>pathname</i></code>	* Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.
Optional arguments	
<code>--force</code>	Forces an existing plan file with the same name to be overwritten.

* - argument is positional.

Description

The following rules apply to the `batch-migrate create-plan` command:

- 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, the command prints a warning to the console, but still creates the plan.

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

Examples

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.

3. Use the `batch-migrate start` command to start the migration:

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

4. 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
```

5. 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 now reside on the target devices.

6. Use `batch-migrate clean` to clean the migration:

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

This dismantles the source devices down to their storage volumes and renames the target devices and volumes using the source device names.

7. Use the `batch-migrate remove` command to remove the record of the migration:

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


Example: Pause/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, 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`
- `batch-migrate check-plan`
- `batch-migrate clean`
- `batch-migrate commit`
- `batch-migrate pause`
- `batch-migrate remove`
- `batch-migrate resume`
- `batch-migrate start`
- `batch-migrate summary`

batch-migrate pause

Pauses the specified batch migration.

Contexts

All contexts.

Syntax

```
batch-migrate pause
```

```
[--file] pathname
```

Arguments

Required arguments	
<code>--file</code> <i>pathname</i>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.

Description

Pauses every migration in the batch. If the command encounters an error, the command prints a warning and continues until every migration has been processed.

You can pause active migrations (a migration that has been started) and resume that migration at a later time.

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

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

- Resume the migration during periods of low I/O.
Use the `batch-migrate resume --file pathname` command to resume the specified paused migration.

Examples

The following example pauses all of the migrations listed in `BSO_19`.

```
VPlexcli:/> batch-migrate pause --file BSO_19
```

See also

- `batch-migrate cancel`
- `batch-migrate check-plan`
- `batch-migrate clean`
- `batch-migrate commit`
- `batch-migrate create-plan`
- `batch-migrate remove`
- `batch-migrate resume`
- `batch-migrate start`
- `batch-migrate summary`

batch-migrate remove

Removes the record of the completed batch migration.

Contexts

All contexts.

Syntax

```
batch-migrate remove  
[--file] Required arguments
```

Arguments

Required arguments	
<code>--file <i>pathname</i></code>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> 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.

Examples

Remove a group of migration jobs.

```
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`
- `batch-migrate check-plan`
- `batch-migrate clean`
- `batch-migrate commit`
- `batch-migrate create-plan`
- `batch-migrate pause`
- `batch-migrate resume`
- `batch-migrate start`
- `batch-migrate summary`

batch-migrate resume

Attempts to resume every migration in the specified batch.

Contexts

All contexts.

Syntax

```
batch-migrate resume
```

```
[--file] pathname
```

Arguments

Required arguments	
<code>--file</code> <i>pathname</i>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> 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.

Examples

Resume all of the migrations specified in the file BSO_19.

```
VPlexcli:/> batch-migrate resume --file BSO_19
```

See also

- `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 start`
- `batch-migrate summary`

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	
<code>--file</code> <i>pathname</i>	* Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.
Optional arguments	
<code>[-s transfer-size]</code> <i>size</i>	<p>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 below. Valid values must be a multiple of 4 K.</p> <ul style="list-style-type: none"> • Range: 40 K - 128 M. • Default: 128 K.

<code>--force</code>	Do not ask for confirmation when starting individual migrations. Allows this command to be run using a non-interactive script, .
<code>--paused</code>	Starts the migration in a paused state. The migration remains paused until restarted using the <code>batch-migrate resume</code> command.

* - argument is positional.

Description

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

CAUTION Inter-cluster migration of volumes 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, the command prints a warning to the console. The command continues until every migration item completes been processing.

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

Batch migrations across clusters can result in the following error:

```

VPlexcli:/> batch-migrate start /var/log/VPlex/cli/migrate.txt
The source device 'SveTest_tgt_r0_case2_1_0002' has a volume
'SveTest_tgt_r0_case2_1_0002_vol' in a view. Migrating to device
'SveTest_src_r0_case2_2_0002' will create a synchronous distributed device.
In this GEO system, this can increase the per I/O latency on
'SveTest_tgt_r0_case2_1_0002_vol'. If applications using
'SveTest_tgt_r0_case2_1_0002_vol' are sensitive to this latency, they may
experience data unavailability. Do you wish to proceed ? (Yes/No) y
WARNING: Failed to start migration /clusters/cluster-1/devices/
SveTest_tgt_r0_case2_1_0002 --> /clusters/cluster-2/devices/
SveTest_src_r0_case2_2_0002 : Evaluation of <<dm migration start -n BR0_0 -f /
clusters/cluster-1/devices/SveTest_tgt_r0_case2_1_0002 -t /clusters/cluster-2/
devices/SveTest_src_r0_case2_2_0002 -s 128kB>> failed.
Failed to create a new data-migration.
Unable to attach mirror 'SveTest_src_r0_case2_2_0002' to distributed Device
'MIGRATE_BR0_0'.
Firmware command error.
Active metadata device does not have a free slot.
Started 0 of 1 migrations.

```

Refer to the troubleshooting section of the VPLEX procedures in the Solve Desktop for instructions on increasing the number of slots.

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 recommended 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.

Examples

```
VPlexcli:/> batch-migrate start --file BS0_19 --transfer-size 1M
```

See also

- `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 summary`
- `dm migration start`

batch-migrate summary

Displays a summary of the batch migration.

Contexts

All contexts.

Syntax

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

Arguments

Required arguments	
<code>--file</code> <i>pathname</i>	Directory and filename of migration plan file. Relative paths can be used. If no directory is specified, the default directory is <code>/var/log/VPlex/cli</code> on the management server.
Optional arguments	

<code>[-v verbose]</code>	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.

Table 5 batch migration summary field descriptions

Field	Description
<code>--verbose</code> output only	
source device	Local-device from which to migrate.
source cluster	Cluster on which source local-device is located.
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.
<code>--verbose</code> and non <code>--verbose</code> output	
Processed n 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.

Table 5 batch migration summary field descriptions (continued)

Field	Description
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 ?.

Examples

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
    
```

Display a batch migration using the --verbose option:

```

VPlexcli:/> batch-migrate summary batch-migrate2.txt --verbose
Command output:
source device      source      target device      target      migration
status  percentage  eta          cluster        -----  cluster        name
-----  -----  ---
-----  -----  ---
-----  -----  ---
temp1_r1_0_cluster-1  cluster-1  temp2_r1_0_cluster-2  cluster-2  BR1_0
complete 100
temp1_r1_1_cluster-1  cluster-1  temp2_r1_1_cluster-2  cluster-2  BR1_1
complete 100
temp1_r1_2_cluster-1  cluster-1  temp2_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
    
```

See also

- 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`

battery-conditioning disable

Disables battery conditioning on the specified backup battery units.

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

Optional arguments	
<code>[-s --sps-unit]</code> <i>context path,context path...</i>	<p>* Standby power supplies (SPS) units on which to disable battery conditioning. If this argument is used:</p> <ul style="list-style-type: none"> • Do not specify the <code>--all-at-cluster</code> argument • The command ignores the <code>--bbu-type</code>.
<code>[-c --all-at-cluster]</code> <i>cluster</i>	<p>* The cluster on which to disable battery conditioning on all backup battery units. If this argument is used:</p> <ul style="list-style-type: none"> • You must specify the <code>--bbu-type</code> argument • Do not use the <code>--sps-units</code> argument
<code>[-t --bbu-type]</code> <i>bbu-type</i>	Type of battery unit on which to disable conditioning. For the current release, only standby-power-supply (SPS) units are supported. The command ignores this argument if you enter the <code>--sps-unit</code> argument.
<code>[-f --force]</code>	Skips the user confirmation that appears when the battery unit on which conditioning is being disabled is currently undergoing conditioning. Allows the command to execute 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.

Examples

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

```
VPlexcli:/> 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'.
```

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

```
VPlexcli:/> 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'.
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                             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
```

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

```
VPlexcli:/> 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`

- `battery-conditioning manual-cycle cancel-request`
- `battery-conditioning manual-cycle request`
- `battery-conditioning set-schedule`
- `battery-conditioning summary`

battery-conditioning enable

Enables conditioning on the specified backup battery units.

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

Optional arguments	
<code>[-s --sps-unit] <i>context path,context path...</i></code>	<p>* Standby power supplies (SPS) units on which to enable battery conditioning. If this argument is used:</p> <ul style="list-style-type: none"> • Do not specify the <code>--all-at-cluster</code> argument • The command ignores the <code>--bbu-type</code> argument
<code>[-c --all-at-cluster] <i>cluster</i></code>	<p>*The cluster on which to enable battery conditioning on all backup battery units. If this argument is used:</p> <ul style="list-style-type: none"> • You must specify the <code>--bbu-type</code> argument • Do not specify the <code>--sps-units</code> argument
<code>[-t --bbu-type] <i>bbu-type</i></code>	<p>Type of battery unit on which to enable conditioning. For the current release, only standby-power-supply (SPS) units are supported. The command ignores this argument if you specify the <code>--sps-unit</code> argument.</p>

* - 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.

Examples

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

```

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
- battery-conditioning manual-cycle cancel-request
- battery-conditioning manual-cycle request
- battery-conditioning set-schedule

- battery-conditioning summary

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	
<code>[-s --sps-unit]</code> <i>context path</i>	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.

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

Examples

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

```
VPlexcli:/> 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
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                           Sat Dec 31 17:00:00 MST 2011
previous-cycle                       Fri Dec 16 21:31:00 MST 2011
previous-cycle-result                 PASS
```

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

```
VPlexcli:/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
- battery-conditioning enable
- battery-conditioning manual-cycle request
- battery-conditioning set-schedule
- battery-conditioning summary

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	
<code>[-s --sps-unit]</code> <i>context path</i>	* Standby power supply (SPS) units 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	
<code>[-f --force]</code>	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 does 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 runs in the next 6-hour window available for the selected unit.

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

Examples

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
```

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`
- `battery-conditioning enable`
- `battery-conditioning manual-cycle cancel-request`
- `battery-conditioning set-schedule`
- `battery-conditioning summary`

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	
<code>[-t --bbu-type] <i>bbu-type</i></code>	* Type of battery backup unit to be conditioned.  Note: In the current release, the only <code>bbu-type</code> supported is <code>sps</code> .
<code>[-d -- day-of-week] [sunday monday ... saturday]</code>	* Day of the week on which to run the battery conditioning. Valid values are: <code>sunday</code> , <code>monday</code> , <code>tuesday</code> , <code>wednesday</code> , etc.
<code>[-c --cluster] <i>cluster</i></code>	* 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.

The following table shows battery conditioning fields:

Table 6 battery conditioning field descriptions

Field	Description
<code>enabled</code>	Whether battery conditioning is enabled on the backup battery unit.
<code>in-progress</code>	Whether a battery conditioning cycle (either manually requested or automatic) is currently underway on the backup battery unit.
<code>manual-cycle-requested</code>	Whether a manually requested battery conditioning cycle is scheduled on the backup battery unit.
<code>next-cycle</code>	The date and time of the next battery conditioning cycle (either manually requested or automatic) on the backup battery unit. 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

Table 6 battery conditioning field descriptions (continued)

Field	Description
	the two time zones, the day selected by the --day-of-week 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 the "next cycle" field above.
previous-cycle-result	<ul style="list-style-type: none"> • ABORTED - The previous battery conditioning cycle was stopped while it was underway. • PASS - The previous battery conditioning cycle was successful. • SKIPPED - The previous cycle was not performed because battery conditioning was disabled. • UNKNOWN - The time of the previous battery conditioning cycle cannot be determined or no cycle has ever run on the backup battery unit.

Examples

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

```

See also

- `battery-conditioning disable`
- `battery-conditioning enable`
- `battery-conditioning manual-cycle cancel-request`
- `battery-conditioning manual-cycle request`
- `battery-conditioning summary`

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`

Description

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

Table 7 battery conditioning summary 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. 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 <code>--day-of-week</code> argument may not be the same day displayed in this field.
Owner	VPLEX engine associated with the specified battery.
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 <code>Next Cycle</code> .

Table 7 battery conditioning summary field descriptions (continued)

Field	Description
Previous Cycle Result	<ul style="list-style-type: none"> ABORTED - The previous battery conditioning cycle was stopped while it was underway. PASS - The previous battery conditioning cycle was successful. SKIPPED - The previous cycle was not performed because battery conditioning was disabled. UNKNOWN - 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 <code>battery-conditioning set-schedule</code> command.
Unit	ID of the power supply where the specified battery is installed.

Examples

Display battery conditioning schedule for all devices:

```
VPlexcli: /> 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

See also

- battery-conditioning disable
- battery-conditioning enable

- battery-conditioning manual-cycle cancel-request
- battery-conditioning manual-cycle request
- battery-conditioning set-schedule

cache-invalidate

In the virtual-volume context, invalidates cached reads on all directors based on its visibility to the clusters. In the consistency-group context, invalidates the cached reads of all exported virtual volumes in the specified consistency group, on all directors based on the virtual volume’s visibility to the clusters.

Contexts

virtual-volume

consistency-group

Syntax

In virtual-volume context:

```
virtual-volume cache-invalidate
    [-v|--virtual-volume] virtual-volume
    [--force]
    [--verbose]
```

In consistency-group context:

```
consistency-group cache-invalidate
    [-g|--consistency-group ] consistency-group context path
    [--force]
    [--verbose]
```

Arguments

Required arguments	
<code>[-v --virtual-volume]</code> <i>virtual-volume</i>	* Invalidates the cache only for the specified virtual volume. Wildcard patterns (CLI glob patterns) are NOT allowed.
<code>[-g --consistency-group]</code> <i>consistency-group context path</i>	* Invalidates the cache for all exported virtual volumes in the consistency-group. Wildcard patterns (CLI glob patterns) are NOT allowed.
Optional arguments	
<code>[--force]</code>	The <code>--force</code> option suppresses the warning message and executes the <code>cache-invalidate</code> directly.
<code>[--verbose]</code>	The <code>--verbose</code> option provides more detailed messages in the command output when specified.

* - argument is positional.

Note: This command does suspend I/O on the virtual volume while invalidation is in progress. Use of this command causes data unavailability on the virtual volume for a short time in a normal operating setup.


Description

In the `virtual-volume` context, invalidates cache reads of a virtual volume on all directors based on its visibility to the clusters. In the `consistency-group` context, invalidates cache reads of all exported virtual volumes in the specified consistency group, on all directors based on the virtual volume's visibility to the clusters.

All cached reads associated with the selected volume are invalidated on all directors of the VPLEX cluster. Subsequent reads from host application will fetch data from storage volume due to cache miss.

Cache can be invalidated only for exported virtual volumes. There are no cache reads on non-exported volumes.

When executed for a given virtual volume, based on the visibility of the virtual volume, this command invalidates the cache reads for that virtual volume on the directors at clusters that export them to hosts.

 **Note:** The CLI must be connected to all directors in order for the cache reads to be flushed and invalidated. The command performs a pre-check to verify all directors are reachable.

When executed for a given consistency group, based on the visibility of each of the virtual volumes in the consistency group, this command invalidates the cache reads for the virtual volumes part of the consistency-group on the directors at clusters that export them to hosts.

The command performs pre-checks to verify all expected directors in each cluster are connected and in healthy state before issuing the command.

If I/O is in progress on the virtual volume, the command issues an error requesting the user to stop all host applications accessing the virtual-volume first.

The command issues a warning stating that the cache invalidate could cause a data unavailability on the volume that it is issued on.

Any cached data at the host or host applications will not be cleared by this command.

Do not execute the `cache-invalidate` command on a RecoverPoint enabled virtual-volume.

The VPLEX clusters should not be undergoing an NDU while this command is being executed.

There should be no mirror rebuilds or migrations in progress on the virtual volume on which the command is executed.

There should be no volume expansion in progress on the virtual volume that the command is being executed.

The applications accessing data on the virtual volume should be stopped. The application or the host operating system should not have any access to the VPLEX virtual volume.

There are four categories of `cache-invalidate` results tracked and reported back by the command:

- If the command completes successfully, the result output will state that the operation was successful. The `cache-invalidate-status` command does not track or report any result for this operation, as the command completed successfully.
- If the command cannot be run because the volume is not exported from any directors. If any of the configured directors are not exporting the virtual volume, then an informational message is displayed and the command output ends.
- If the command failed on a director. The result displays an error indicating that the command failed on a specific director.
- If the command execution exceeds five minutes and there are no directors with a failed response, a check is made on directors that have the operation running in background. The result displays an error that shows the directors with the operation still running, and an instruction to execute the `cache-invalidate-status` command to view progress.

Note: Results from the `virtual-volume cache-invalidate` and `consistency-group cache-invalidate` commands will be basically the same with different target types. Different messages will be returned for the default brief mode messages and for verbose mode messages when the `--verbose` option is specified.

Examples

Example output when cache invalidation for the virtual volume and consistency group is successful.

Brief message:

```
Cache invalidation successful for virtual-volume 'testvol_1'.
Cache invalidation successful for consistency-group 'test_cg_1'.
```

Verbose message:

```
Cache invalidation was successful.
Cache Invalidation Summary  Volume Names
-----
invalidate completed      testvol_1
                           testvol_2
not registered with cache  testvol_3 consistency-group 'test_cg_1'.
```

Example output when cache-invalidation failed for the virtual volume and for the consistency group.

Brief message:

```
Cache invalidation failed for consistency-group 'test_cg_1'.
Cache invalidation failed for virtual-volume 'test_volume_1'.
```

Verbose message:

```
Cache invalidation failed.
Cache Invalidation Summary  Director Names
-----
I/O error                  director-1-1-A
                           director-1-1-B
                           director-2-1-A
invalidate completed       director-2-1-B
```

Example output when any directors cannot be reached. The message includes an instruction to re-issue the command.

Note: The `cache-invalidate` command must be completed for every director where the virtual volume or consistency group is potentially registered in cache.

Virtual-volume:

```
Cache invalidation for virtual-volume 'testvol_1' could not be performed on
directors 'director1-1-A, director-1-1-B'.
Please make sure all configured directors are connected and healthy and re-
issue the command.
```

Consistency group:

```
Cache invalidation for consistency-group 'test_cg_1' could not be performed
on directors 'director1-1-A, director-1-2-A'.
Please make sure all configured directors are connected and healthy and re-
issue the command.
```

Example output returned when a `cache-invalidate` command for a virtual volume or consistency group has timed out after five minutes.

Virtual volume:

```
Cache invalidation for virtual-volume 'test_volume_1' is not complete on the
following directors: director-1-1-A. Please use the 'virtual-volume cache-
invalidate-status test_volume_1' command to find out the current state of
cache invalidation.
```

Consistency group:

```
Cache invalidation for consistency-group 'test_cg_1' is not complete on the
following directors: director-1-1-A. Please use the 'consistency-group cache-
invalidate-status test_cg_1' command to find out the current state of cache
invalidation.
```

See also

- `cache-invalidate-status`
- `virtual-volume destroy`
- `virtual-volume provision`

cache-invalidate-status

Displays the current status of a `cache-invalidate` operation on a virtual volume, or for a specified consistency group.

Contexts

`virtual-volume`

`consistency-group`

Syntax

In virtual-volume context:

```
cache-invalidate-status
[-v|--virtual-volume] virtual-volume
[-h|--help]
[--verbose]
```

In consistency-group context:

```
cache-invalidate-status
[-g|--consistency-group] consistency-group context path
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-v --virtual-volume] <i>virtual-volume</i></code>	* Specifies the virtual volume cache invalidation status to display.
<code>[-g --consistency-group] <i>consistency-group context path</i></code>	* Specifies the consistency group cache invalidation status to display.
Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides a detailed summary of results for all directors when specified. If not specified, only the default brief summary message is displayed.

* - argument is positional.

Description

This command is executed only when the `virtual-volume cache-invalidate` command or the `consistency-group cache-invalidate` command has exceeded the timeout period of five minutes. If this command is executed for a reason other than timeout, the command will return “no result.”

The `cache-invalidate status` command displays the current status for operations that are successful, have failed, or are still in progress.

If the `--verbose` option is not specified, this command generates a brief message summarizing the current status of the cache-invalidation for the virtual volume or consistency group, whichever context has been specified. In the event of cache-invalidation failure, the brief mode output will only summarize the first director to fail, for if one director fails, the entire operation is considered to have failed.

If the `--verbose` option is specified, then the `cache-invalidate status` is displayed on each of the directors on which the `cache-invalidate` is pending.

There are three fields of particular interest

Table 8 Important cache invalidate status fields

Field	Description
status	<p>indicates the status of the cache-invalidate.</p> <ul style="list-style-type: none"> • <code>in-progress cache-invalidate</code> is running. • <code>completed cache-invalidate</code> is finished. • <code>error cache-invalidate</code> is in error state. • <code>unknown</code> Unable to retrieve the information. This could be because the director is not reachable. The reason will be reflected in the <code>cause</code> field. • – Nothing is there to report.

Table 8 Important cache invalidate status fields (continued)

result	<p>Indicates the result of the cache-invalidate.</p> <ul style="list-style-type: none"> • successful cache-invalidate finished successfully. • failure cache-invalidate finished with failure. The reason for the failure is reflected in the cause field. • error cache-invalidate is in an error state. • – Nothing to report.
cause	<p>Displays the information for the failure of a job command. <code>text</code> An error message explaining the reason for failure.</p>

This command is read-only. It does not make any changes to storage in VPLEX or to VPLEX in any way. It is only for providing information on the current status of cache-invalidation operations.

Examples

Example output when the cache-invalidate operation finished with success.

Brief message:

```
Cache invalidation for consistency-group 'test_cg_1' is complete.
```

Verbose message:

```
VPlexcli:/> virtual-volume cache-invalidate-status vv1
cache-invalidate-status
-----
director-1-1-A      status:  completed
                   result:  successful
                   cause:  -
director-2-1-B      status:  completed
                   result:  successful
                   cause:  -
VPlexcli:/> consistency-group cache-invalidate-status test_cg_1
cache-invalidate-status
-----
director-1-1-A      status:  completed
                   result:  successful
                   cause:  -
director-1-1-B      status:  completed
                   result:  successful
                   cause:  -
```

Example output when the cache-invalidate operation finished with failure.

Brief message:

```
Cache invalidation for virtual-volume 'vv1' failed on director 'director-1-1-A' because: Volume is in suspended state.
Cache invalidation for consistency-group 'test_cg_1' failed on director 'director-1-1-A' because: I/O error.
```

Verbose message:

```

VPlexcli:/> virtual-volume cache-invalidate-status vv1
cache-invalidate-status
-----
director-1-1-A      status:  completed
                   result:  failure
                   cause:  Volume is in suspended state
VPlexcli:/> consistency-group cache-invalidate-status test_cg_1
cache-invalidate-status
-----
director-1-1-A      status:  completed
                   result:  failed
                   cause:  I/O error
director-1-1-B      status:  completed
                   result:  failed
                   cause:  I/O error
    
```

Example output when the cache-invalidate operation is still in progress.

Brief message:

```

Cache invalidation for virtual-volume 'vv1' is still in-progress.
Cache invalidation for consistency-group 'test_cg_1' is still in-progress.
    
```

Verbose message:

```

cache-invalidate-status
-----
director-1-1-A      status:  in-progress
                   result:  -
                   cause:  -
director-1-1-B      status:  in-progress
                   result:  -
                   cause:  -
    
```

See also

- `cache-invalidate`

capture begin

Begins a capture session.

Contexts

All contexts.

Syntax

```

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

Arguments

Required arguments	
--------------------	--

<code>[-s --session] <i>session name</i></code>	* Name of capture session. Output files from the capture session are named using this value.
<code>[-c --capture-directory] <i>directory</i></code>	* Pathname for the capture directory. Default capture directory: <code>/var/log/VPlex/cli/capture</code>

* - 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 commands issued during the capture session.
- `session name-stdin.txt` - CLI commands input during the capture session.
- `session name-stdout.txt` - Output of commands 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.

Examples

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`
- `capture pause`
- `capture replay`
- `capture resume`

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.

Examples

End a capture session.

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

See also

- capture begin
- capture pause
- capture replay
- capture resume

capture pause

Pauses the current capture session.

Contexts

All contexts.

Syntax

```
capture pause
```

Description

Pause/resume operates only on the current capture session.

Examples

Pause a capture session.

```
VPlexcli:/> capture pause
```

See also

- capture begin
- capture end
- capture replay
- capture resume

capture replay

Replays a previously captured session.

Contexts

All contexts.

Syntax

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

Arguments

Required arguments	
<code>[-s --session] <i>session name</i></code>	* Name of existing capture session.
<code>[-c --capture-directory] <i>directory</i></code>	* Directory where existing captured session is located. Default directory <code>/var/log/VPlex/cli/capture/recapture</code>

* - 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

Replay a capture session.

```
VPlexcli:/> capture replay TestCapture
Attributes:
Name                               Value
-----
allow-auto-join                     true
auto-expel-count                    0
auto-expel-period                   0
.
.
.
```

See also

- capture begin
- capture end
- capture pause
- capture resume

capture resume

Resumes the current capture session.

Contexts

All contexts.

Syntax

```
capture resume
```

Description

Pause/resume operates only on the current capture session.

Examples

Resume the current capture session.

```
VPlexcli:/> capture resume
```

See also

- capture begin
- capture end
- capture pause
- capture replay

cd

Changes the working directory.

Contexts

All contexts.

Syntax

```
cd [context]
```

Arguments

Optional arguments	
<i>context</i>	<p>Change to the specified context. The context can be one of the following:</p> <ul style="list-style-type: none"> • context path - The full or relative pathname of the context. • .. - the parent context of the context you are currently in. • ... - the root context. • -(dash) - The context you were in before changing to this context. <p>If you do not specify a <i>context</i>, the <code>cd</code> command changes to the root directory.</p>

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*.

Examples

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	
--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 series column	The column in the CSV file to use as series.
--range series range	The column in the CSV file to use as range.
--domain domain column -	The column in the CSV file to use as domain.
--width chart width	The width of chart graphic. <ul style="list-style-type: none"> • Range: 64-2048. • Default: 500.
--height chart height	The height of the chart graphic. <ul style="list-style-type: none"> • Range: 64-2048. • Default: 500.
--aggregate aggregate-series-name	Name of aggregate series. If not specified, no aggregate series is displayed.

All arguments are positional.

Examples

In the following example:

- The `exit` command exits the CLI and returns to the management server.
- The `cd` command navigates to the `reports` context
- 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` context
- 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 volumes, 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>
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`
- `report capacity-arrays`
- `report capacity-clusters`
- `report capacity-hosts`

- report create-monitors

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	
<code>[-c --cluster] <i>context path</i></code>	* Cluster to add.
<code>[-t --to] <i>cluster</i></code>	* Cluster to which the given cluster is added. This is only necessary if the system cannot be automatically determined.
Optional arguments	
<code>[-f --force]</code>	Forces the cluster addition to proceed even if conditions are not optimal.

* - argument is positional.

Description

Before a cluster can communicate with the other cluster of a Metro, you must use the cluster add command.

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.

Examples

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	TLA	Connected	Expelled	Operational	Status	Health
cluster-1	1	FNM00103600160	true	false	ok	ok	
cluster-2	2	FNM00103600161	true	false	ok	ok	
Islands:							
Island ID	Clusters						
1	cluster-1, cluster-2						

See also

- cluster expel
- cluster status
- cluster summary

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	
<code>[-c --clusters] <i>clusters,cluster</i></code>	Flushes the cache for every exported virtual volume of every director at the specified clusters. Entered as wildcard patterns.
<code>[-v --volumes] <i>volumes</i></code>	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	
<code>[-e --sequential]</code>	Flushes the cache of multiple directors sequentially. Default is to flush the caches in parallel.
<code>--verbose</code>	Displays a progress report during the flush. Default is to display no output if the run is successful.

Description

Note: 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.

Examples

```

VPlexcli:/clusters/> cacheflush --clusters /clusters/* --verbose --sequential
Flushing director1 director2 ...
  Flushing logging_logging1_vol @ director1... done
  Flushing regression_dR1a_l2_vol @ director1... done
  Flushing regression_dR1b_l2_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_l2_vol @ director2... done
  Flushing regression_dR1b_l2_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
- cluster summary

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

Optional arguments	
<code>[-c --clusters] <i>cluster</i></code>	Dump configuration information for only the specified cluster.
<code>[-d --dtdOnly]</code>	Print only the Document Type Definitions (DTD) document.

<code>[-f --file] <i>filename</i></code>	Direct the configdump output to the specified file. Default location for the output file on the management server is: <code>/var/log/VPlex/cli</code> .
--	---

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.

Examples

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

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

Dump the configuration at cluster-1, navigate to the cli context 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>
  <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`
- `director appcon`
- `getsysinfo`

- sms dump

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	
<code>[-c --clusters] <i>cluster</i></code>	* The cluster to expel.
<code>[-f --force]</code>	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, 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.

Examples

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  TLA                Connected  Expelled  Operational
  Status   Health State
  -----
cluster-1  1          FNM00103600160    true       true
isolated  degraded
```

```

cluster-2 2          FNM00103600161 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

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] <i>context path</i>	* 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.

Examples

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

See also

- cluster add
- cluster expel
- cluster status
- cluster unexpel

cluster restart-local-cluster

Initiates cluster restart and restores local cluster configuration on VPLEX VS2 Metro configuration.

Context

All

Syntax

```
cluster restart-local-cluster
-h | --help
--verbose
-p | --pre-check
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.
-p --pre-check	Performs cluster pre-restart validation.

Description

`cluster restart-local-cluster` supports cluster restart on VS2 Metro configuration. Follow manual cluster shutdown and restart procedure available in the SolVe Desktop for cluster restart scenarios:

Run this command on the same cluster that was shutdown using `cluster stop-local-cluster`.

This command is cluster specific. Run this command on the management server of the cluster that was shutdown. For example, if you shutdown cluster-1, run the command `cluster restart-local-cluster` the Cluster-1 management server.

The management server and cluster-witness IP addresses should remain the same before restarting the cluster.

cluster show-remote-devices

Displays the list of remote devices for the specified cluster.

Contexts

All contexts.

Syntax

`cluster show-remote-devices options cluster`

Description

The command displays the list of remote devices for the specified cluster. The top-level volumes and the list of views at which the devices are exported are also listed. Use the `--verbose` option to see the complete list.

Arguments

Required arguments	
options	
positional arguments	
Optional arguments	
<code>-h --help</code>	Displays the usage for the command.
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.
<code>-s --include-sub-devices</code>	Displays all remote RAIDs. If sub devices are not specified, the command displays only the top-level RAIDs.
<code>-[-c --cluster=] cluster context</code>	Specifies the context path of the cluster to show the remote devices.

Examples

```
VPlexcli:/> cluster show-remote-devices -c cluster-1 --include-sub-devices
```


Remote Device	Virtual Volume	Cluster	Views
c2_dev	c2_dev_vol	cluster-1	[View_67]
c2_Dr_device0049_2	-	-	-
c2_Dr_device0048_2	-	-	-
c2_Dr_device0047_2	vol2	-	-
c2_Dr_device0046_2	vol3	cluster-1	[view1, view2]
c2_Dr_device0045_2	-	-	-
c2_Dr_device0044_2	-	-	-
c2_Dr_device0043_2	vol4	cluster-2	[view3]
c2_Dr_device0042_2	-	-	-
c2_Dr_device0041_2	-	-	-

(181 more)
To see all results please run the command with --verbose option.

Show top-level devices only:

```
VPlexcli:/> cluster show-remote-devices -c cluster-2
```

Remote Device	Virtual Volume	Cluster	Views
device_remotevol1_c1	remotevol1	cluster-2	[View_49]

Show remote devices with verbose option:

```
VPlexcli:/> cluster show-remote-devices -c cluster-1 --verbose
```

Remote Device	Virtual Volume	Cluster	Views
c2_dev	c2_dev_vol	cluster-1	[View_67]
c2_Dr_device0049_2	-	-	-
c2_Dr_device0048_2	-	-	-
c2_Dr_device0047_2	vol2	-	-
c2_Dr_device0046_2	vol3	cluster-1	[view1, view2]
c2_Dr_device0045_2	-	-	-
c2_Dr_device0044_2	-	-	-
c2_Dr_device0043_2	vol4	cluster-2	[view3]
c2_Dr_device0042_2	-	-	-
c2_Dr_device0041_2	-	-	-

```

c2_Dr_device0040_2 - - -
c2_Dr_device0039_2 - - -
c2_Dr_device0038_2 - - -
c2_Dr_device0037_2 - - -
c2_Dr_device0036_2 - - -
c2_Dr_device0035_2 - - -
c2_Dr_device0034_2 - - -
c2_Dr_device0033_2 - - -

```

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	
<code>[-c --cluster] <i>context path</i></code>	Cluster to shut down.
<code>[-f --force]</code>	Forces the shutdown to proceed.

Description

 **WARNING** Shutting down a VPLEX cluster could cause data unavailability. Please refer to the VPLEX procedures in the SolVe Desktop for the recommended procedure to shut down a cluster.

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 procedures in the SolVe Desktop 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)'
```

Examples

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.

```
VPlexcli:/> cluster shutdown -c cluster-1
VPlexcli:/> cluster shutdown 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. To show that
you understand the impact enter 'SHUTDOWN': SHUTDOWN

You have chosen to shutdown 'cluster-1'. To confirm, enter 'cluster-1': cluster-1

Status      Description
-----
Started.    Shutdown started.
VPlexcli:/> cluster summary
Clusters:
  Name      Cluster ID TLA      Connected  Expelled
  Operational Status Health State
  -----
  cluster-1 1          FNM00103600160  true       false
  unknown
  cluster-2 2          FNM00103600161  true       false
ok
Islands:
  Island ID Clusters
  -----
  1         cluster-1, cluster-2
VPlexcli:/> cluster summary
Clusters:
Name      Cluster ID TLA      Connected
Expelled Operational Status Health State
-----
cluster-1 1          FNM00103600160  false      -
-
cluster-2 2          FNM00103600161  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
- cluster expel
- cluster forget
- director shutdown

cluster status

Displays a cluster's operational status and health state.

Contexts

All contexts.

Syntax

```
cluster status
```

Description

The following table shows the fields displayed in the `cluster status` command:

Table 9 cluster status field descriptions

Field	Description
<code>operational status</code>	<p>Operational status of the cluster. During transition periods cluster moves from one operational state to another.</p> <ul style="list-style-type: none"> • <code>cluster departure</code> - One or more of the clusters cannot be contacted. Commands affecting distributed storage are refused. • <code>degraded</code> - The location is not functioning at an optimal level. This may indicate non-functioning remote virtual volume, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, out-of-date devices, and so forth. • <code>device initializing</code> - If clusters cannot communicate with each other, then the distributed-device will be unable to initialize.

Table 9 cluster status field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>device out of date</code> - Child devices are being marked fully out of date. Sometimes this occurs after a link outage. • <code>expelled</code> - The cluster has been isolated from the island either manually (by an administrator) or automatically (by a system configuration setting). • <code>ok</code> - The cluster is operating normally. • <code>shutdown</code> - The cluster's directors are shutting down. • <code>suspended exports</code> - 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. It may not be a problem, and the VPLEX might be waiting for you to confirm the resumption of I/O. • <code>transitioning</code> - 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).
<code>transitioning - indications</code>	Additional information if the <code>transitioning-progress</code> is anything other than blank.
<code>transitioning-progress</code>	Indicates progress for supported transitions.
<code>health-state</code>	<p><code>critical failure</code> - The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity.</p> <ul style="list-style-type: none"> • <code>degraded</code> - The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volume, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, or out-of-date devices. • <code>ok</code> - The cluster is functioning normally. • <code>unknown</code> - VPLEX cannot determine the cluster's health state, or the state is invalid. • <code>major failure</code> - The cluster is failing and some functionality may be degraded or unavailable. This may indicate complete loss of back-end connectivity. • <code>minor failure</code> - The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.
<code>health-indications</code>	Additional information if the <code>health-state</code> field is anything other than <code>ok</code> .
<code>local-com</code>	<code>ok</code> - All wan-com links have the expected connectivity: this port-group is operating correctly.

Table 9 cluster status field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>warning</code> - Some links have unexpected connectivity. This port-group is operational but not properly configured. Performance may not be optimal. • <code>error</code> - Some connectivity is missing from this port-group. It is not operating correctly. • <code>fail</code> - All connectivity is missing from this port-group. <code>wan-com</code> is not operational.
<code>Certificates</code>	<ul style="list-style-type: none"> • <code>ok</code> - The system has a valid certificate. • <code>warning</code> - The certificate will expire in 30 days. Expired certificates can impact inter-cluster connectivity, which can lead to data unavailability. Renew the certificates using the <code>security renew-all-certificates</code> command before their expiry. • <code>error</code> - The certificate has expired, or a certificate is not found on the cluster. Expired certificates can impact inter-cluster connectivity and it can lead to data unavailability. If a certificate is expired, renew it using the <code>security renew-all-certificates</code> command. If a certificate is not found on the system, create a new certificates using the <code>security configure-certificates</code> command.
<code>wan-com</code>	<p><code>full</code> - All port-groups have a status of either <code>ok</code> or <code>warning</code>. <code>wan-com</code> connectivity is complete through minor configuration errors may still exist. See individual port-group statuses.</p> <ul style="list-style-type: none"> • <code>partial</code> - Some port-groups have a status of <code>error</code> or <code>fail</code>, but at least one port-group has a status of <code>ok</code> or <code>warning</code>. <code>wan-com</code> is operating (possibly minimally) through at least one channel. Performance is degraded. • <code>none</code> - All port-groups have a status of either <code>error</code> or <code>fail</code>. <code>wan-com</code> is not operational. • <code>not-applicable</code> - The system is a single-cluster (i.e. Local) system. Validating <code>wan-com</code> connectivity is not applicable.
<code>license</code>	Whether the license is installed or not.

See also

- `cluster summary`
- `ds summary`

cluster stop-local-cluster

Perform graceful shutdown of a single local cluster in VPLEX VS2 Metro configuration.

Context

All

Syntax

```
cluster stop-local-cluster
  [-h | --help]
  [--verbose]
  [-p | --pre-check]
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.
-p --pre-check	Performs cluster pre-shutdown preparation and validation.

Description

`cluster stop-local-cluster` supports cluster shutdown on a VS2 Metro configuration. Follow the manual cluster shutdown procedures available in the SolVe Desktop for these cluster shutdown scenarios:

- Shutdown cluster in VPLEX local configuration (All VS1, VS2 and VS6 platforms)
- Shutdown both the clusters in VPLEX Metro configuration (All VS1, VS2 and VS6 platforms)
- Shutdown single cluster in VS6 VPLEX Metro configuration

Note:

- This command is cluster specific. Run this command on the management server of the cluster to be shutdown. For example, to shut down Cluster-1, run the command `cluster stop-local-cluster` on the cluster-1 management server only.
- Remote volumes should be manually migrated as the size of remote volumes may be very huge and user should have disk space to migrate these volumes.
- Cluster shutdown on MetroPoint configuration is not supported.

3.

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`

Description

The following table shows the fields available in the `cluster summary` output.

Table 10 cluster summary field descriptions

Field	Description
Clusters:	
Name	Name of the cluster.
Cluster ID	For VPLEX Local, always 1. For VPLEX Metro, 1 or 2.
TLA	The Top-level Assembly. The product TLA must uniquely identify the product instance. For VPLEX the TLA must uniquely identify the cluster (which is the rack and all physical components in it)
Connected	Whether or not the CLI is connected to at least one director in the cluster (connected to the cluster). <ul style="list-style-type: none"> • <code>true</code> - CLI is connected to the cluster. • <code>false</code> - CLI is not connected to the cluster.
Expelled	<ul style="list-style-type: none"> • <code>true</code> - The cluster is expelled from its island. • <code>false</code> - The cluster is not expelled from its island.
Operational Status	<ul style="list-style-type: none"> • <code>degraded</code> - 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 hardware health problems, suspended virtual volumes, suspended exports, storage volumes not visible from all directors, meta-volume not yet processed. • <code>isolated</code> - The cluster is not communicating with any other clusters. • <code>ok</code> - The cluster is functioning normally. • <code>transitioning</code> - The cluster is reacting to external events and may not be operating as configured. I/O may be suspended during the transition period.

Table 10 cluster summary field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> If no meta-volume has been configured, operational status is <code>transitioning</code>. <code>unknown</code> - 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	<ul style="list-style-type: none"> <code>critical failure</code> - The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity. <code>degraded</code> - 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. <code>ok</code> - The cluster is functioning normally. <code>unknown</code> - The VPLEX cannot determine the cluster's health state, or the state is invalid. <code>major failure</code> - The cluster is failing and some functionality may be degraded or unavailable. This may indicate complete loss of back-end connectivity. <code>minor failure</code> - The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.
Islands:	
Island ID	ID of the island. For current release, always 1.
Clusters	Names of clusters belonging to the island. For current release, always <code>cluster-1</code> or <code>cluster-2</code> .

Examples

Display summary for healthy clusters:

```
Vplexcli: /> cluster summary
Clusters:
```

```

Name      Cluster ID  TLA      Connected  Expelled  Operational
Status   Health State
-----
cluster-1 1          FNM00190701072 true      false
ok
cluster-2 2          FNM00190701073 true      false
ok

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

```

Display cluster summary for VPLEX Metro configuration with a inter-cluster link outage:

```

Vplexcli:/> cluster summary
Clusters:
Name      Cluster ID  TLA      Connected  Expelled  Operational
Status   Health State
-----
cluster-1 1          FNM00190701072 true      false
ok
cluster-2 2          FNM00190701073 true      false
ok
degraded
degraded

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

```

Display cluster summary for VPLEX Metro configuration with a cluster expelled:

```

Vplexcli:/> cluster summary
Clusters:
Name      Cluster ID  TLA      Connected  Expelled  Operational
Status   Health State
-----
cluster-1 1          FNM00190701072 true      true
isolated degraded
cluster-2 2          FNM00190701073 true      true
isolated degraded

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

```

See also

- cluster status

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	
<code>[-c --cluster] <i>context path</i></code>	Cluster to unexpel.

Description

Clears the expelled flag for the specified cluster, allowing it to rejoin the VPLEX.

Examples

To manually unexpel a cluster, do the following:

1. Use the `cluster summary` command to verify that the cluster is expelled.

```
VPlexcli:/> cluster summary
Clusters:
  Name          Cluster ID  TLA          Connected  Expelled  Operational
  Status Health State
  -----
cluster-1  1          ok          FNM00190701072  true      false
ok
cluster-2  2          ok          FNM00190701073  true      false
ok

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

2. Use the `ll` command in the target cluster's `cluster` context to display the cluster's `allow-auto-join` attribute setting.

```
VPlexcli:/> ll /clusters/cluster-1
/custers/cluster-1:
Attributes:
Name          Value
-----
allow-auto-join  true
auto-expel-count  0
auto-expel-period  0
auto-join-delay  0
cluster-id      1
.
```

```
.
.
```

If the cluster's `allow-auto-join` attribute is set to `true`, the cluster automatically rejoins the system. Skip to step 4.

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-1
VPlexcli:/clusters/cluster-1> set allow-auto-join true
```

4. Use the `cluster unexpel` command to manually unexpel a cluster, allowing the cluster to rejoin VPlex. The syntax for the command is:
For example:

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

5. Use the `cluster summary` command to verify all clusters are in one island and working as expected.

```
VPlexcli:/>cluster summary
Clusters:
  Name          Cluster ID  TLA          Connected  Expelled  Operational
  Status Health State
  -----
  cluster-1  1          FNM00091300128 true       false
ok
  cluster-2  2          FNM00091300218 true       false
ok
Islands:
  Island ID  Clusters
  -----
  1          cluster-1, cluster-2
```

See also

- `cluster expel`

cluster-witness configure

Creates the cluster-witness context for enabling VPlex Witness functionality and configuration commands.

Contexts

All contexts.

Syntax

```
cluster-witness configure
[--verbose]
```

Arguments

Optional arguments	
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Cluster Witness is an optional component of VPLEX Metro 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.

Note: This command must be run on both management servers to create cluster-witness CLI contexts on the VPLEX.

The following must be true for the command to run successfully:

- The Cluster Witness must be disabled.
 - The VPN from this management server to the Cluster Witness server must be established and functional.
 - The Cluster Witness server must be operational and connected.
- Note:** 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.

Examples

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/          recoverpoint/
security/
VPlexcli:/> cluster-witness configure
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

- Dell EMC VPLEX Procedures in SolVe Desktop - "VPLEX Witness: Install and Setup"
- `cluster-witness disable`
- `cluster-witness enable`
- `configuration cw-vpn-configure`

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`.

Syntax

```
cluster-witness disable
[-f|--force]
[-w|--force-without-server]
```

Arguments

i **Note:** This command is available only after Cluster Witness has been configured and `cluster-witness` CLI context is visible.

Optional arguments	
<code>[-f --force]</code>	Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.
<code>[-w --force-without-server]</code>	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 <code>--force-without-server</code> 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. i Note: If Cluster Witness Server becomes reachable when <code>--force-without-server</code> option is used, the command will also disable the Cluster Witness Server.

Description

Disables Cluster Witness on both management servers and on the Cluster Witness server.

Allows consistency group rule-sets to dictate I/O behavior to distributed virtual volumes in consistency groups.

Note: 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 meta-volume on both clusters is healthy

Examples

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"
```

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

- Procedures in SolVe Desktop - "VPLEX Witness: Install and Setup"
- `cluster summary`
- `cluster-witness enable`
- `vpn status`
- `cluster-witness configure`
- `configuration cw-vpn-configure`

cluster-witness enable

Enables Cluster Witness on both clusters and on the Cluster Witness Server.

Contexts


All contexts.


In `/cluster-witness` context, command is enable.

Syntax

```
cluster-witness enable
```

Description


 **Note:** This command is available only after Cluster Witness has been configured and `cluster-witness` CLI context is visible.

 **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 system 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 meta-volume 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 section of the SolVe Desktop and/or contact Dell EMC Customer Support.

The `cluster-witness` context does not appear in the 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`

Table 11 cluster witness display fields

Field	Description
<code>Name</code>	Name of component, <ul style="list-style-type: none"> • For cluster – name assigned to cluster. • For VPLEX Witness server – “server”.
<code>id</code>	ID of a VPLEX cluster. Always blank “-“ for Witness server.
<code>admin state</code>	Identifies whether VPLEX Witness is enabled/disabled. Valid values are: <ul style="list-style-type: none"> • <code>enabled</code> - VPLEX Witness functionality is enabled on this component. • <code>disabled</code> - VPLEX Witness functionality is disabled on this component. • <code>inconsistent</code> - 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. • <code>unknown</code> - This component is not reachable and its administrative state cannot be determined.
<code>private-ip-address</code>	Private IP address of the Cluster Witness server virtual machine used for Cluster Witness-specific traffic.
<code>public-ip-address</code>	Public IP address of the Cluster Witness server virtual machine used as an endpoint of the IPsec tunnel.
<code>diagnostic</code>	String generated by CLI based on the analysis of the data and state information reported by the corresponding component. <pre>WARNING: Cannot establish connectivity with Cluster Witness Server to query diagnostic information. - Cluster Witness Server or one of the clusters is unreachable. Local cluster-x hasn't yet established connectivity with the server - The cluster has never connected to Cluster Witness</pre>

Table 11 cluster witness display fields (continued)

	<p>Server.</p> <p>Remote cluster-x hasn't yet established connectivity with the server - The cluster has never connected to Cluster Witness Server.</p> <p>Cluster-x has been out of touch from the server for X days, Y secs - Cluster Witness Server has not received messages from a given cluster for longer than 60 seconds.</p> <p>Cluster witness server has been out of touch for X days, Y secs - Either cluster has not received messages from Cluster Witness Server for longer than 60 seconds.</p> <p>Cluster Witness is not enabled on component-X, so no diagnostic information is available - Cluster Witness Server or either of the clusters is disabled.</p>
<p>Operational State</p>	<p>Operational state of the corresponding component.</p> <ul style="list-style-type: none"> • For clusters – Healthy state is <code>in-contact</code>. All other states indicate a problem. • For Cluster Witness server – Healthy state is <code>clusters-in-contact</code>. All other states indicate a problem. <p>For Cluster Witness server:</p> <ul style="list-style-type: none"> • <code>clusters-in-contact</code> - Both clusters are in contact with each other over the inter-cluster network. • <code>cluster-partition</code> - 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. • <code>cluster-unreachable</code> - 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. • <code>unknown</code> - 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. • When the server's operational state is <code>cluster-partition</code> or <code>cluster-unreachable</code>, 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

Table 11 cluster witness display fields (continued)

	<p>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> <p>For cluster:</p> <ul style="list-style-type: none"> • <code>in-contact</code> - 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 cluster in all synchronous consistency groups is allowed from the perspective of VPLEX Witness. • <code>cluster-partition</code> - 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 cluster in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of VPLEX Witness. • <code>remote-cluster-isolated-or-dead</code> - 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 cluster in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of VPLEX Witness. • <code>local-cluster-isolated</code> - 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 cluster in all synchronous consistency groups regardless of cluster preference. <ul style="list-style-type: none"> ⓘ 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 <code>local-cluster-isolated</code> if this cluster remains partitioned from the peer cluster. • <code>unknown</code> - 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. <ul style="list-style-type: none"> ⓘ Note: When its state is <code>local-cluster-isolated</code>, the cluster does not receive guidance from Cluster Witness server. If connectivity with Cluster Witness server is
--	--

Table 11 cluster witness display fields (continued)

	restored before the inter-cluster link is restored state is local-cluster-isolated.
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> <ul style="list-style-type: none"> ok - The component is reachable failed - The component is not reachable

Examples

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 VPLEX 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

```

See also

- Procedures in the SolVe Desktop
- cluster summary
- cluster-witness disable
- vpn status
- cluster-witness configure
- configuration cw-vpn-configure

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

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: <code>getsysinfo</code> , <code>export port summary</code> , and <code>connectivity director</code> .
<code>--local-only</code>	Gathers diagnostics only from the local cluster and directors.
<code>--minimum</code>	Combines all space-saving and time-saving options and operates only on the local cluster. Use this argument when the <code>collect-diagnostics</code> command is expected to take a long time or produce excessively large output files. Combines the <code>--noextended</code> , <code>--faster</code> , and <code>--local-only</code> arguments. Also omits the second <code>sms dump</code> command and output file (<code>smsDump-CLI_Logs_timestamp.zip</code>).
<code>--allcores</code>	Collect all available core files from the directors. By default, only the two latest core files are collected.
<code>--large-config</code>	Omits the cluster <code>configdump</code> command output. Use this argument only when the configuration has: <ul style="list-style-type: none"> • 8000 or more storage volumes, and • 4000 or more local top-level devices, and • 2000 or more distributed devices
<code>--recoverpoint-only</code>	Collects only RecoverPoint diagnostic information. If there are no cores files to be collected, the command generates only one base file.
<code>--out-dir <i>directory</i></code>	The directory into which to save the zip file containing the collected diagnostics. Default: <code>/diag/collect-diagnostics-out</code> Not needed for normal usage.
<code>-r --log-range=<i>Range</i></code>	Collect only the logs generated in the range specified. To collect the set of logs generated in a certain range of hours of the current day, specify the value as <code>start-hour-end-hour</code> . For example, to collect the logs between 11AM to 1PM, enter the <code>range 11-13h</code> . To collect the set of logs generated in a certain range of days in the current month, enter the range <code>start-day-end-day</code> d. For example, to get the logs between 11 days ago to 13 days ago, enter the <code>range of 11-13d</code> .
<code>-l --last-logs=<i>x</i></code>	Collect only the logs generated in the last <i>x</i> days or hours are collected. To collect the logs generated in the last 4 hours, the value for this option is 4h. To collect the logs generated in last 4 days, use the value 4d.

Description

Collects logs, cores, and configuration information from the management server and directors. Places the collected output files in the `/diag/collect-diagnostics-out` directory on the management server.

Two compressed files are placed in the `/diag/collect-diagnostics-out` directory:

- `tla-diagnostics-extended-timestamp.tar.gz` - Contains java heap dump, fast trace dump, two latest core files, and two latest RecoverPoint kdriver core files (if they exist).
- `tla-diagnostics-timestamp.tar.gz` - Contains everything else including a directory, `/opt/recoverpoint` which contains the RecoverPoint splitter logs in a zip file (`vpsplitter.log.xx`, `vpsplitter.log.periodic_env`, and `vpsplitter.log.current_env`).

Best practice is to collect both files. The extended file is usually large, and thus takes some time to transfer.

Recommended practice is to transfer the base file (`tla-diagnostics-timestamp.tar.gz`) first and begin analysis while waiting for the extended file to transfer.

Note: On VPLEX Metro 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 unless the `--local-only` argument is used.

All trace files under the folder `/cores/nsfw-trace/` on each director will be copied and put in the collected diagnostics.

Note: Core files are always collected only from local directors. 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 configurations, run the `collect-diagnostics` command on each management server, but NOT at the same time. Even if you use the `--local-only` argument, do not run the command on both management servers at the same time.

If you know the time period in which an event happened, you can collect a subset of logs based on time period. This reduces the time to collect diagnostics. Use the `--log-range` and `--last-logs` arguments to collect a subset of diagnostics.

Remove files created by `collect-diagnostics` from the management server as soon as possible to avoid filling management server disk partitions.

Examples

Collect diagnostics, omitting the core files on the directors and the management server console heap, and send the output to the default directory:

```
Vplexcli:/> collect-diagnostics --nocores --noheap
```

Collect all RecoverPoint diagnostics, including all available RecoverPoint core files, and send the output to the default directory:

```
Vplexcli:/> collect-diagnostics --recoverpoint-only --allcores
```

See also

- `cluster configdump`
- `director appdump`
- `getsysinfo`
- `sms dump`

configuration complete-system-setup

Completes the configuration.

Contexts

All contexts.

Syntax

```
configuration complete-system-setup
```

Description

Completes the automated EZ-Setup Wizard.

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 cluster).

Examples

```
VPlexcli:/> configuration complete-system-setup
```

See also

- `configuration connect-remote-directors`
- `configuration continue-system-setup`
- `configuration sync-time-clear`
- `configuration system-setup`

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 the `authentication directory-service configure` command for a description of the available authentication services.

Examples

Configure the selected authentication service:

```
VPlexcli:/> 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 mapprincipal []: 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`
- `authentication directory-service unconfigure`

configuration connect-local-directors

Connects to the directors in the local cluster.

Contexts

All contexts.

Syntax

```

configuration connect-local-directors
[-f|--force]

```

Arguments

Optional arguments	
<code>[-f --force]</code>	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`.

Examples

Connect the local directors to the cluster:

```

VPlexcli:/> configuration connect-local-directors

```

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 connect-remote-directors
- configuration complete-system-setup
- configuration system-setup

configuration connect-remote-directors

Connects to the remote directors after the VPN connection has been established.

Contexts

All contexts.

Syntax

```
configuration connect-remote-directors
[-c|--engine-count] engine count
[-f|--force]
```

Arguments

Optional arguments	
<code>[-c --engine-count] <i>engine count</i></code>	Specifies the number of engines present at the remote site.
<code>[-f --force]</code>	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.

Examples

Connect remote directors to the directors in the local cluster:

```
VPlexcli:/> configuration connect-remote-directors --engine-count 2
Connected to Plex firmware director-2-1-A.
Connected to Plex firmware director-2-1-B.
Connected to Plex firmware director-2-2-A.
Connected to Plex firmware director-2-2-B.
```

See also

- `configuration connect-local-directors`
- `configuration complete-system-setup`
- `configuration system-setup`

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 you complete the back-end storage configuration and allocation for the cluster, use this command to complete the initial configuration.

Examples

```
Vplexcli:/> configuration continue-system-setup
```

See also

- `configuration system-setup`

configuration cw-vpn-configure

Establishes VPN connectivity between a VPLEX management server and the Cluster Witness Server and starts the VPN tunnel between them.

Contexts

All contexts.

Syntax

```
configuration cw-vpn-configure
[-i|--ip-address] public-ip-address
[-c | --cwsHostCert=] host-certificate
-k | --cwsHostKey= hostkey
-f | --force
-h | --help
--verbose
```

Arguments

Required arguments	
<code>[-i --ip-address] <i>public-ip-address</i></code>	* Valid public IP address of the Cluster Witness server.
Optional arguments	
<code>-h --help</code>	Displays the usage for this command.
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.
<code>[-c --cwsHostCert=] <i>host-certificate</i></code>	Specifies the absolute path of the Cluster Witness server host certificate to import.
<code>-k --cwsHostKey= <i>hostkey</i></code>	Specifies the absolute path of Cluster Witness server host key name to import
<code>-f --force</code>	Force the update of the configuration file of the Cluster Witness server.

Description

The command is interactive and requires inputs to complete successfully.

Note: This command must be run on both management server, precisely on cluster-1 followed by cluster-2 to establish the VPN between the management server and the Cluster Witness server.

A management server authenticates an external client entity based on the Certificate Authority it trusts. The trust store is used for web/ REST clients over https connections and ssl database for inter-site VPN connections. The CA trust can be self-signed based on a local CA subject info or a third party vendor (such as Verisign, Globalsign) signed. Executing the command without `-c` and `-k` options creates the self-signed Cluster Witness server host certificate. signed by the same CA which is used in establishing the VPN between the two management server. To import the Cluster Witness server host certificate and key, run the command with `--cwsHostCert` or `-c` and `--cwsHostKey` and `-k` options respectively.

Note: The Cluster Witness server certificate is to be signed by the same CA which has signed the certificates that are used in establishing the VPN between the two management server. In the case of self-sign, this is taken care by the command itself. In the case of import, please provide the Cluster Witness server certificate signed by the same CA which has signed the certificates that are used in establishing the VPN between the two management server.

Prerequisites

Before executing the command ensure the following conditions are met:

- VPLEX Metro setup is successfully completed using EZSetup. This creates a VPN connection between VPLEX management servers.
- The Cluster Witness server certificate is to be signed by the same CA which has signed the certificates that are used in establishing the VPN between the two management server.

Note: If this command is run when the VPN is already configured, the following error message is displayed: `VPN connectivity is already established.`

configuration cw-change-password

This command updates the Cluster Witness Server password.

Contexts

All contexts.

Syntax

```
configuration cw-change-password
[-p | --prompt]
[-c | --promptCurrentPassword]
[-f | --force]
-h | --help
--verbose
```

Arguments

Optional arguments	
<code>[-p --prompt]</code>	Prompts for a new Cluster Witness Server password.
<code>[-c --promptCurrentPassword]</code>	Prompts for current password instead of retrieving current password from the lockbox. This option is typically used if the Cluster Witness Server password has been previously changed.
<code>[-f --force]</code>	Forces Cluster Witness Server password update. This option is used to change the current password, if the password is not the default password
<code>-h --help</code>	Displays usage for this command.
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

This command is used to change the Cluster Witness Server password.

The Cluster Witness Server password is identical to the management server password by default. It is imperative that the Cluster Witness password be changed to provide system security.

The *Dell EMC VPLEX Security Configuration Guide* provides recommendations about choosing an appropriately complex password.

If a customer attempts to change a password when the password is not the default password, the following informational message is displayed: `Cluster Witness default password is already changed.`

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 configurations.

This command should be used only when Cluster Witness is disabled, because it is not providing guidance to VPLEX clusters at that time.

Using the `cw-vpn-reset` command when Cluster Witness is enabled causes Cluster Witness to lose connectivity with the management server and generates a call-home event.

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.

Examples

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
- cluster-witness disable
- cluster-witness enable

configuration enable-front-end-ports

After the meta-volume is created using the EZ-Setup wizard, enable front-end ports using this command.

Contexts

All contexts.

Syntax

```
configuration enable-front-end-ports
```

Description

Completes the initial system configuration using the EZ-Setup Wizard. After configuring the meta-volume 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.

Examples

Enable the front-end ports.

```
VPlexcli: /> configuration enable-front-end-ports
```

See also

- meta-volume create
- configuration metadata-backup
- configuration complete-system-setup

- `configuration system-setup`

configuration event-notices-reports config

Configure call-home notification and system reporting 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 call-home notification.

If call-home notification is already configured, the current configuration information is displayed.

If call-home notification 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 notification:

- IP address of the primary SMTP server used to forward reports and call-home notifications to Dell EMC. Dell EMC recommends using your ESRS gateway as the primary connection address.
- (Optional) One or more IP address of secondary SMTP servers used to forward reports and call-home notifications to Dell 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 messages when call-home notifications occur.

See also

- `configuration event-notices-reports reset`
- `configuration event-notices-reports show`
- `notifications call-home test`
- `scheduleSYR list`

configuration event-notices-reports reset

Resets the current call-home notification and system reporting configuration.

Contexts

All contexts.

Syntax

```
configuration event-notices-reports-reset
```


Description

This command:

- Deletes the current call-home notification and system reporting configuration data
- Disables call-home notification

See also

- `configuration event-notices-reports config`
- `configuration event-notices-reports show`
- `notifications call-home test`
- `scheduleSYR list`

configuration event-notices-reports-show

This command shows call-home notification connection records and system connections based on the ConnectEMC configuration.

Contexts

All contexts.

Syntax

```
configuration event-notices-reports-show
```

```
[-h | --help]
```

```
[--verbose]
```

Arguments

Optional arguments	
<code>-h --help</code>	Displays the usage for this command.
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

This command shows call-home notification connection records and system connections based on the ConnectEMC configuration. There are no limitations.

Connection types are:

Table 12 Connection Types

Primary	The first SMTP address to be used to send notifications to Dell EMC.
Failover	Alternate SMTP address to send notifications to Dell EMC. Failover connections will be attempted, in order, when the primary and any previous failover attempts have failed.
onSuccess/onFailure	Sends an email to the customer each time Dell EMC is successfully sent an email notification.

Table 12 Connection Types (continued)

Primary	The first SMTP address to be used to send notifications to Dell EMC.
onFailure	Sends an email to the customer each time an attempt to notify Dell EMC has failed.
onAllFailure	Sends an email to the customer only if all attempts to notify Dell EMC have failed.
onSuccess	Sends an email to the customer regardless of whether the email notification to Dell EMC succeeded.

See also

- `configuration event-notices-reports config`
- `configuration event-notices-reports show`

configuration flashdir-backup disable

Disables periodic director flashdir backups.

Contexts

All contexts.

Syntax

```
configuration flashdir-backup disable
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

Disables periodic director flashdir backups.

Example

Disable flashdir backups.

```
VPlexcli:/> configuration flashdir-backup disable
Flashdir backups have been disabled
```

See also

- `configuration flashdir-backup enable`

configuration flashdir-backup enable

Enables periodic director flashdir backups.

Contexts

All contexts.

Syntax

```
configuration flashdir-backup enable
```

```
[-h|--help]
```

```
[--verbose]
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Enables periodic director flashdir backups.

Example

enable backup of flashdir.

```
VPlexcli:/> configuration flashdir-backup enable
Flashdir backups have been disabled
```

See also

- `configuration flasjdir-backup disable`

configuration get-product-type

Displays the VPLEX product type (Local or Metro).

Contexts

All contexts.

Syntax

```
configuration get-product-type
```

Description

Displays whether the system is a Local or Metro configuration .

Example

Display the configuration type.

```

Vplexcli:/> configuration get-product-type
The cluster is currently configured as a VPLEX Metro

```

See also

- cluster status
- cluster summary
- version

configuration join-clusters

Validates WAN connectivity and joins the two clusters.

Contexts

All contexts.

Syntax

```

configuration join-clusters
[-i|--remote-ip] remote IP address
[-h|--help]

```

Arguments

Optional arguments	
<code>[-i --remote-ip] <i>remote IP address</i></code>	Specifies the IP address of the remote server.
<code>[-h --help]</code>	Displays command line help.

Description

This command validates WAN connectivity and joins the two clusters.

 **Note:** This command can be configured as Metro Fibre Channel using the EZ-Setup wizard.

Example

Join clusters at the specified remote IP address:

```

Vplexcli:/> configuration join-clusters -i 10.103.97.76
Verifying the connectivity of all the directors...
Verifying that all the pre-conditions for director connectivity are
satisfied...
All the directors are properly connected to each other...
Verifying island ID
Running the Join Cluster task
Verifying island ID
Added cluster 'cluster-1' to system (cluster-2).
Join Cluster task completed.
The clusters are now joined and ready for use.

```

See also

- cluster add
- configuration continue-system-setup
- configuration system-setup

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.

The meta-volume's backup size should be equal to or greater than the active meta-volume size. The current requirement is 78G per storage volume.


See the Dell EMC VPLEX Technical Notes for best practices regarding the kind of back-end array volumes to consider for a meta-volume.

 **Note:** This command must be executed on the management server in which you want to create the backups.

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** Renaming backup metadata volumes is *not* supported. 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,VPD83T3:600009700001926017145330364
64237
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]:

```

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 remote-clusters clear-addresses
- configuration show-meta-volume-candidates

- `configuration system-setup`

configuration register-product

Registers the VPLEX product with Dell EMC.

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.

The 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 Dell EMC. Two methods are available:

- Attach the registration file to an e-mail and send it to Dell EMC.
- Send the registration file to Dell EMC through an SMTP server.
If this option is selected, a prompt for an SMTP server IP address is displayed.

Examples

Register the product.

```
Vplexcli:/> 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
- configuration system-setup

configuration remote-clusters add-addresses

Adds entries to the /clusters/*/connectivity/wan-com::remote-cluster-addresses list.

Context

All contexts.

Syntax

```

configuration remote-clusters add-addresses
[-c |--cluster] cluster context
[-d |--default]
[-r |--remote-cluster] cluster context
[-a |--addresses] addresses[, addresses...]
[-h |--help]
[--verbose]
    
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster context</i></code>	* Specifies the cluster whose connectivity configuration is to be modified. Typically the cluster above the current context.
<code>[-d --default]</code>	Applies default configuration values. Default configuration pulls the cluster-address attribute of the active subnets from all remote clusters. This option does not require <code>--remote-cluster</code> or <code>--addresses</code> .

<code>[-r --remote-cluster] <i>cluster context</i></code>	Specifies the remote-cluster configuration entry to modify. Cannot be the same context specified in <code>--cluster</code> . This option is not required when <code>--default</code> is used. <code>--addresses</code> is required if this option is used.
<code>[-a --addresses] <i>addresses [, addresses...]</i></code>	Specifies one or more remote <i>ip-address:port-number</i> entries to add for the specified <code>--remote-cluster</code> . This option is not required when <code>--default</code> is used. If this option is used <code>--remote-cluster</code> is also required.
<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Adds one or more address:port configurations for the specified remote-cluster entry for this cluster.

See the VPLEX procedures in the SolVe Desktop for more information on managing WAN-COM IP addresses.

Examples

Add a cluster 2 address to cluster 1:

```
VPlexcli:/> ls /clusters/cluster-1/connectivity/wan-com
/clusters/cluster-1/connectivity/wan-com:
Attributes:
Name                               Value
-----
discovery-address                   224.100.100.100
discovery-port                       10000
listening-port                       11000
remote-cluster-addresses             cluster-2 [192.168.11.252:11000]
VPlexcli:/> configuration remote-clusters add-addresses -c cluster-1/ -r cluster-2/ -a
10.6.11.252:11000
VPlexcli:/> ls /clusters/cluster-1/connectivity/wan-com
/clusters/cluster-1/connectivity/wan-com:
Attributes:
Name                               Value
-----
discovery-address                   224.100.100.100
discovery-port                       10000
listening-port                       11000
remote-cluster-addresses             cluster-2 [10.6.11.252:11000, 192.168.11.252:11000]
```

Add cluster 2 addresses to cluster 1, assuming that the WAN COM port groups for cluster 2 have their cluster-address properly configured:

```
VPlexcli:/clusters/cluster-1/connectivity/wan-com> ls
Attributes:
Name                               Value
-----
discovery-address                   224.100.100.100
discovery-port                       10000
listening-port                       11000
remote-cluster-addresses             -
VPlexcli:/clusters/cluster-1/connectivity/wan-com> remote-clusters add-addresses -d
VPlexcli:/clusters/cluster-1/connectivity/wan-com> ls
```

Attributes:	
Name	Value
discovery-address	224.100.100.100
discovery-port	10000
listening-port	11000
remote-cluster-addresses	cluster-2 [10.6.11.252:11000, 192.168.11.252:11000]

See also

- configuration remote-clusters clear-addresses

configuration remote-clusters clear-addresses

Removes entries from the /clusters/*/connectivity/wan-com::remote-cluster-addresses list.

Context

All contexts.

Syntax

```
configuration remote-clusters clear-addresses
[-c | --cluster] cluster context
[-r | --remote-cluster] cluster context
[-a | --addresses] addresses[,addresses]
[-h | --help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster context</i></code>	* Specifies the cluster whose connectivity configuration is to be modified. Typically the cluster above the current context.
<code>[-r --remote-cluster] <i>cluster context</i></code>	* Specifies the remote-cluster configuration entry to modify. Cannot be the same context specified in --cluster.
<code>[-a --addresses] <i>addresses [, addresses...]</i></code>	Specifies one or more remote <i>ip-address:port-number</i> entries to remove for the specified --remote-cluster. If no entry is specified, all entries are removed for the specified --remote-cluster.
<code>[-h --help]</code>	Displays the usage for this command.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Clears one, several, or all address:port configurations for the specified remote-cluster entry for this cluster.

See the VPLEX procedures in the SolVe Desktop for more information on managing WAN-COM IP addresses.

Examples

Clear a cluster 2 address on cluster 1:

```

VPlexcli:> ls /clusters/cluster-1/connectivity/wan-com
/clusters/cluster-1/connectivity/wan-com:
Attributes:
Name                               Value
-----
discovery-address                   224.100.100.100
discovery-port                       10000
listening-port                       11000
remote-cluster-addresses            cluster-2 [10.6.11.252:11000, 192.168.11.252:11000]
VPlexcli:> configuration remote-clusters clear-addresses -c cluster-1/-r cluster-2/-a
10.6.11.252:11000
VPlexcli:> ls /clusters/cluster-1/connectivity/wan-com
/clusters/cluster-1/connectivity/wan-com:
Attributes:
Name                               Value
-----
discovery-address                   224.100.100.100
discovery-port                       10000
listening-port                       11000
remote-cluster-addresses            cluster-2 [192.168.11.252:11000]

```

See also

- `configuration remote-clusters add-addresses`

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 you configure the meta volume on a CLARiiON® array, do not configure the meta volume on the vault drives of the CLARiiON.

Dell 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. Isolate the physical spindles for meta volumes 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.

Examples

Show meta volume candidates:

```

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

- meta-volume create
- configuration metadata-backup
- configuration system-setup

configuration subnet clear

Clears values of all attributes in one or more subnet contexts, with exception of maximum transfer unit attribute (mtu).

Contexts

```

/ : configuration subnet clear
/clusters/*/connectivity/*/port-groups/ip-port-group-* : subnet clear
/clusters/*/connectivity/*/port-groups/ip-port-group-*/subnet/ : clear

```

Syntax

```
clear [--force] subnet-context[,subnet-context[, ...]]
```

Description

If a list of subnet contexts is specified, the subnets are cleared.

You can also specify non-subnet contexts. In this case, the command clears all subnet contexts in ip-port-groups below the specified context (as if `*/subnet` were appended to the given path).

This command destroys configuration information, thereby requiring user confirmation before executing. The list of subnets to be cleared is presented for confirmation. You can bypass this check by specifying the `--force` option.

The configuration `subnet clear` command works for all subnet contexts under `/clusters/*/connectivity/`.

Subnet properties cannot be modified unless both of the following are true:

- Ports in the port-group are disabled
- Directors in the local cluster are reachable

Examples

Clear all subnets in `ip-port-group-3`:

```
VPlexcli:/clusters/cluster-1/connectivity/wan-com/port-groups/ip-port-group-3/
subnet> ls
Name                               Value
-----
cluster-address                    192.168.10.10
gateway                            192.168.10.1
mtu                                1500
prefix                             192.168.10.0/255.255.255.0
proxy-external-address             10.10.42.100
remote-subnet-address              192.168.100.10/255.255.255.0
VPlexcli:/clusters/cluster-1/connectivity/wan-com/port-groups/ip-port-group-3/
subnet> clear ./
Clearing subnet for the following port-groups:
  Cluster   Role   Port-Group
  -----
  cluster-1 wan-com ip-port-group-3
Are you sure you want to clear these subnets? (Yes/No) yes
VPlexcli:/clusters/cluster-1/connectivity/wan-com/port-groups/ip-port-group-3/
subnet> ls
Name                               Value
-----
cluster-address                    -
gateway                            -
mtu                                1500
prefix                             -
proxy-external-address             -
remote-subnet-address              -
```

Clearing all subnets of `ip-port-group-3` explicitly:

```
VPlexcli:/> ls /clusters/cluster-1/connectivity/wan-com/port-groups/ip-port-
group-3/subnet
Name                               Value
-----
cluster-address                    192.168.10.10
gateway                            192.168.10.1
mtu                                1500
prefix                             192.168.10.0/255.255.255.0
proxy-external-address             10.10.42.100
remote-subnet-address              192.168.100.10/255.255.255.0
VPlexcli:/> configuration subnet clear /clusters/cluster-1/connectivity/wan-
com/port-groups/ip-port-group-3/subnet/
Clearing subnet for the following port-groups:
  Cluster   Role   Port-Group
  -----
  cluster-1 wan-com ip-port-group-3
Are you sure you want to clear these subnets? (Yes/No) yes
```

Clearing all subnets of cluster-1 with confirmation:

```

VPlexcli:/> configuration subnet clear clusters/cluster-1/
Clearing subnet for the following port-groups:
Cluster      Role      Port-Group
-----
cluster-1    wan-com   ip-port-group-4
cluster-1    wan-com   ip-port-group-3
cluster-1    local-com ip-port-group-1
cluster-1    local-com ip-port-group-2
cluster-1    front-end ip-port-group-2
cluster-1    front-end ip-port-group-1
cluster-1    back-end  ip-port-group-2
cluster-1    back-end  ip-port-group-1
Are you sure you want to clear these subnets? (Yes/No) yes
    
```

configuration subnet remote-subnet add

Adds a routing entry for a remote subnet to the list of remote-subnets in an iSCSI port-group's subnet.

Contexts

All contexts.

Syntax

```

configuration subnet remote-subnet add
[-s|--subnet] subnet context
[-r|--remote-subnets] prefix[, prefix...]
[-h|--help]
[--verbose]
    
```


Arguments

Required arguments	
<code>[-s --subnet] <i>subnet context</i></code>	* Specifies the front-end or back-end iSCSI subnet where a routing prefix should be added.
<code>[-r --remote-subnets] <i>prefix[, prefix...]</i></code>	* Specifies the routing prefixes (remote subnets) to add to the subnet.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Adds to a front-end or back-end iSCSI subnet, a routing prefix (remote subnet) that should be accessed through the subnet's gateway.

 **Note:** This command is valid only on systems that support iSCSI devices.

Examples

Adding two prefixes to the iSCSI port-group 1 subnet:

```
VPlexcli:/> configuration subnet remote-subnet add /clusters/cluster-1/
connectivity/front-end/port-groups/iscsi-port-group-1/subnet/
192.168.1.0/255.255.255.0,192.168.2.0/255.255.255.0
Successfully added 2 remote subnets.
```

Adding a prefix from within an appropriate subnet context:

```
VPlexcli:/clusters/cluster-1/connectivity/back-end/port-groups/iscsi-port-
group-8/subnet>
Successfully added 1 remote subnets.
```

Attempt to add a prefix in a front-end iSCSI port-group subnet failed:

```
VPlexcli:/> configuration subnet remote-subnet add /clusters/cluster-1/
connectivity/front-end/port-groups/iscsi-port-group-1/subnet/
192.168.1.0/255.255.255.0
There were errors adding remote subnets:
Prefix          Error
-----
192.168.1.0/255.255.255.0 Failed to add entry to remote-subnets on
director-1-1-A
Successfully added 0 remote subnets.
```

See also

- configuration subnet clear
- configuration subnet remote-subnet remove

configuration subnet remote-subnet remove

Removes one or more routing entries for remote subnets from the list of remote-subnets in an iSCSI port-group's subnet.

Contexts

All contexts.

Syntax

```
configuration subnet remote-subnet remove
[-s|--subnet] subnet context
[-r|--remote-subnets] prefix[, prefix...]
[-h|--help]
[--verbose]
```

Arguments


Required arguments	
<code>[-s --subnet] <i>subnet context</i></code>	* Specifies the front-end or back-end iSCSI subnet where a routing prefix should be removed.

<code>[-r --remote-subnets] prefix [, prefix...]</code>	* Specifies the routing prefixes (remote subnets) to remove from the subnet.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Removes from a front-end or back-end iSCSI subnet, a routing prefix (remote subnet) that should be accessed through the subnet's gateway.

 **Note:** This command is valid only on systems that support iSCSI devices.

Examples

Removing two prefixes from the iSCSI port-group 1 subnet:

```
VPlexcli:/> configuration subnet remote-subnet remove /clusters/cluster-1/
connectivity/front-end/port-groups/iscsi-port-group-1/subnet/
192.168.1.0/255.255.255.0,192.168.2.0/255.255.255.0
Successfully removed 2 remote subnets.
```

Removing a prefix from within an appropriate subnet context:

```
VPlexcli:/clusters/cluster-1/connectivity/back-end/port-groups/iscsi-port-
group-8/subnet> remote-subnets remove ./ 192.168.1.0/255.255.255.0
Successfully removed 1 remote subnets.
```

Attempt to remove a prefix from an iSCSI port-group subnet failed:

```
VPlexcli:/> configuration subnet remote-subnet remove /clusters/cluster-1/
connectivity/front-end/port-groups/iscsi-port-group-1/subnet/
192.168.1.0/255.255.255.0
There were errors removing remote subnets:
Prefix          Error
-----
-----
192.168.1.0/255.255.255.0 Failed to add entry to remote-subnets on
director-1-1-A
Successfully removed 0 remote subnets.
```

See also

- configuration subnet clear
- configuration subnet remote-subnet add

configuration sync-time

Synchronizes the time of the local management server with a remote NTP server. Remote Metro clusters are synchronized with the local management server.

Contexts

All contexts.

Syntax


```
configuration sync-time
[-i|--remote-ip] remote-server-IP-address
[-f|--force]
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-i --remote-ip] <i>remote-server-IP-address</i></code>	Specifies the IP address of the remote NTP server.
<code>[-f --force]</code>	Skips the verification prompt.
<code>[-h --help]</code>	Displays the command line help for this command.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

In a VPLEX Metro configuration, the `configuration sync-time` command 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.

This command synchronizes Cluster 1 with the external NTP server, or synchronizes Cluster 2 with the local management server on Cluster 1.

Use this command before performing any set-up on the second cluster of a VPLEX Metro configuration.

Use this command during initial system configuration before using the `configuration system-setup` command.

If you do not provide the IP address of the NTP server or the local management server, the command prompts for the public IP address of the NTP server, or the local management server.

Examples

Running the synchronization time task on cluster-2:

```
Vplexcli:/> configuration sync-time
Running the Sync Time task
Please enter the IP address of the NTP server in order to synchronize the
system clock. ATTENTION: Vplex only supports syncing Cluster-1 to an external
NTP server while Cluster-2 must be synced to Cluster-1. IP Address of NTP
```

```

server:10.108.69.121
Syncing time on the management server of a Metro/Geo system could cause the
VPN to require a restart. Please Confirm (Yes: continue, No: exit) (Yes/No)
yes
Shutting down network time protocol daemon (NTPD)..done
PING 10.108.69.121 (10.108.69.121) 56(84) bytes of data.
64 bytes from 10.108.69.121: icmp_seq=1 ttl=64 time=0.166 ms
--- 10.108.69.121 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.166/0.166/0.166/0.000 ms
Starting network time protocol daemon (NTPD)..done
executing sudo /etc/init.d/ntp stop
executing /bin/ping 10.108.69.121 -c 1
executing sudo /usr/sbin/sntp -r -P no -u 10.108.69.121
executing sudo /etc/init.d/ntp start
Now running 'vpn status' in case the VPN was affected:
Verifying the VPN status between the management servers...
IPSEC is UP
Remote Management Server at IP Address 10.108.69.121 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
Synchronize Time task completed.

```

Running the synchronization task on cluster-1:

```

VPlexcli:/> configuration sync-time --remote-ip 128.222.208.98 -f
Running the Sync Time task
Shutting down network time protocol daemon (NTPD)..done
PING 128.222.208.98 (128.222.208.98) 56(84) bytes of data.
64 bytes from 128.222.208.98: icmp_seq=1 ttl=62 time=0.522 ms
--- 128.222.208.98 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.522/0.522/0.522/0.000 ms
Starting network time protocol daemon (NTPD)..done
executing sudo /etc/init.d/ntp stop
executing /bin/ping 128.222.208.98 -c 1
executing sudo /usr/sbin/sntp -r -P no -u 128.222.208.98
executing sudo /etc/init.d/ntp start
Now running 'vpn status' in case the VPN was affected:
Synchronize Time task completed.

```

See also

- `configuration sync-time-clear`
- `configuration sync-time-show`
- `configuration system-setup`

configuration sync-time-clear

Clears the NTP configuration on a management server.

Contexts

All contexts.

Syntax

```

configuration sync-time-clear
[-h|--help]

```

[--verbose]

Arguments

Optional arguments	
[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

Clears the NTP configuration and restarts the NTP protocol.

This command clears the NTP configuration on the management server of the cluster on which it is run.

Examples

Clears the NTP configuration:

```
VPlexcli:/> configuration sync-time-clear
Shutting down network time protocol daemon (NTPD)..done
Starting network time protocol daemon (NTPD)..done
executing sudo /etc/init.d/ntp restart
```

See also

- configuration sync-time
- configuration sync-time-show
- configuration system-setup

configuration sync-time-show

Displays the NTP configuration of a cluster's management server.

Contexts

All contexts.

Syntax

```
configuration sync-time-show
```

```
[-h|--help]
```

```
[--verbose]
```

Arguments

Optional arguments	
[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

Used to check the management server's NTP configuration.

Examples

Show that no NTP is configured:

```
VPlexcli:/> configuration sync-time-show
No NTP server is configured.
```

Show a server with NTP configured:

```
VPlexcli:/> configuration sync-time-show
This cluster is configured to get its time from the following servers:
10.6.210.192
```

See also

- `configuration sync-time`
- `configuration sync-time-clear`
- `configuration system-setup`

configuration system-reset

Resets the cluster to the manufacturing state.

Contexts

All contexts.

Syntax

```
configuration system-reset
```

Description

Resets the cluster to the manufacturing state.

Use this command to cancel all the configuration completed using the EZ-Setup wizard and return the system to its factory default settings. Any values specified during the configuration session become the defaults displayed when EZ-Setup is re-run.

Run the `configuration system-reset` command only on a new cluster.

Do not run this command on a configured cluster.

For this command to execute, no meta-volume may exist.

Examples

Reset the cluster configuration.

```
VPlexcli:/> configuration system-reset
```

See also

- `configuration system-setup`

configuration system-setup

Starts the EZ-Setup Wizard automated configuration tool.

Contexts

All contexts.

Syntax

```
configuration system-setup
```

Description

Configures the VPN and establishes a secure connection between the clusters

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.

Note: To run this command, there must be no meta-volume configured, and no storage exposed to hosts.

EZ-Setup Wizard automates the following steps:

- Connects to the local directors
- Sets the cluster ID
- Sets the Cluster IP Seed
- Sets the director count
- Commissions the directors
- Forces a service password change
- Configures NTP
- Configures Fibre Channel switch
- Enables COM and back-end ports
- Creates Web server and VPN certificate
- Configures VPN
- Configures/enables call-home
- Configures/enables System Reporting (SYR)

Note: Once the service password is changed, communicate the new password to Dell EMC service and to the VPLEX System Administrator.

Examples

```
Vplexcli: /> configuration system-setup
```

See also

- *Dell EMC VPLEX Configuration Guide*
- *Dell EMC VPLEX Security Configuration Guide*
- About cluster IP seed and cluster ID section in the `security ipsec-configure` command
- `configuration continue-system-setup`

- configuration sync-time-clear

configuration upgrade-meta-slot-count

Upgrades the slot count of the active meta volume at the given cluster to 64,000 slots.

Context

/clusters/cluster/system-volumes

Syntax

```
configuration upgrade-meta-slot-count
```

```
[-c | --cluster=] cluster
```

```
[-d | --storage-volumes=volume [volume, ...]
```

```
[-h | --help ]
```

```
[--verbose]
```

```
[-f | --force]
```

Arguments

Optional arguments	
[-h --help]	Displays the usage for this command.
[--verbose]	Provides more output during command execution.
[-c --cluster=] <i>cluster</i>	The cluster at which to upgrade the slot count of the active meta volume. When specified from within a /clusters/cluster context, the value of that context is used as cluster. The -c or --cluster argument is positional.
[-d --storage-volumes= <i>volume</i> [, <i>volume</i> ...]	Creates a temporary meta volume from one or more storage volumes. After the command completes successfully, the command destroys the temporary 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. Specify two or more storage volumes. Storage volumes should be on different arrays.
[-f --force]	Forces the upgrade to proceed without asking for confirmation.

Description

On the metadata volume, each slot stores header information for each storage volume, extent, and logging volume. This command upgrades the slot count of the active meta volume at the given cluster to 64,000 slots.

By default, the oldest meta volume backup at the cluster serves as a temporary meta volume.

If you specify the -d or --storage-volume option, then the command creates the temporary meta volume from scratch from those disks. The temporary meta volume is active while the currently-active meta volume is being upgraded. At the end of the process, VPLEX reactivates the original meta volume and the temporary meta volume becomes a backup again. VPLEX renames the backup to reflect the new point in time at which it became a backup.

Meta-volumes differ from standard storage volumes in the following ways:

- Create a meta volume from a volum that is not yet claimed.
- Create meta volumes directly on storage volumes, not extents.
If you configure the meta-volume on a CLARiiON array, do not place the meta volume on the vault drives of the CLARiiON.
- Performance is not critical for meta-volumes. The minimum performance allowed is 40MB/sec and 1004KIOP/second.
- Isolate the physical spindles for meta-volumes from application workloads.

Dell EMC recommends the following for meta-volumes:

- Enable read caching
- Pre-configure a hot spare meta-volume in case of a catastrophic failure of the active meta-volume.

If you specify two or more storage-volumes, configure them on two separate arrays if more than one array is present. This command creates a RAID-1 of all the storage-volumes.

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

Optional arguments	
<code>[-o --host] {<i>host-name</i> <i>IP address</i>}</code>	* Host name or IP address of the director to which to connect. Default: localhost.
<code>--logport <i>port-number</i></code>	For use by Dell EMC personnel only. A firmware log event port. Applicable only to test versions of the firmware.
<code>--secondary-host {<i>host-name</i> <i>IP address</i>}</code>	Host name or IP address of the redundant interface on director to which to connect.
<code>--secondary-logport <i>secondary-port-number</i></code>	For use by Dell EMC personnel only. A firmware log event port. Applicable only to test versions of the firmware.

<code>--name <i>name</i></code>	* Name to assign to the director after the connection. Default: Host name or IP address.
<code>[-t --type] <i>type</i></code>	For use by Dell 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.
<code>[-p --password] <i>password</i></code>	Set the password for the connection.
<code>[-c --connection-file] <i>filename</i></code>	Load a list of connections from file named <code>connections</code> at: <code>service@ManagementServer:/var/log/VPlex/cli</code> on the VPLEX management server.
<code>[-s --save-authentication]</code>	Save the authentication information used to connect to the VPLEX without asking for confirmation.
<code>--no-prompt</code>	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 if the management server is not able to connect to 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` 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 host name 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
```

Examples

```
Vplexcli:> connect --host 128.221.252.67 --secondary-host 128.221.253.67 --
name director-2-1-A
```

See also

- disconnect

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|wnn|port]
```

Arguments

Required arguments	
<i>director</i>	
Optional arguments	Director to discover.
[-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] <i>filename</i>	Save the output in the specified file. Default: /var/log/Vplex/cli
[-s --sort-by] {name wnn port}	Sort output by one of the following: <ul style="list-style-type: none"> • name - Sort output by storage volume name. • wnn - Sort output by WorldWide name.

	<ul style="list-style-type: none"> • <code>port</code> - 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.

See also

- `connectivity show`
- `connectivity validate-be`

connectivity list all

Displays the initiators, storage volumes, directors, and the targets connected to a director.

Contexts

```
/: connectivity list all
/engines/*/directors/*/: list all
```

Syntax

```
connectivity list all
[-h | --help]
[--verbose]
[-n | --directors] context path, context path...
[-f | ] filename
[-d | ]
[-i | ]
[-v | ]
[-t | ]
[-s | --sort-by] key
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command
<code> [--verbose]</code>	Provides additional output during command execution. This may not have any effect for some commands.
<code>[-n --directors] <i>context path, context path ...</i></code>	Source director(s) for which connectivity should be reported.
<code>[-f --file] <i>filename</i></code>	Writes the connectivity report to the named file instead of echoing it. If the file exists, any previous contents will be lost.
<code>[-d --show-directors]</code>	Shows inter-director connectivity.
<code>[-i --show-initiators]</code>	Shows connected initiators.

<code>[-v --show-storage-volumes]</code>	Shows connected storage volumes.
<code>[-t --show-targets]</code>	Shows connected targets.
<code>[-s --sort-by] <i>key</i></code>	The field by which to sort the storage volume information (name, wwn or port).

Description

Use this command to list the directors, initiators, storage volumes, and the targets that are connected to a director. Reports the results of the `connectivity list storage-volumes`, `connectivity list initiators`, and the `connectivity list directors` commands for each specified director. Unless you specify `-d`, `-i`, or `-v`, all three categories are reported. The reports are ordered by director, not by report category.

See also

- `connectivity list directors`
- `connectivity list initiators`
- `connectivity list storage-volumes`

connectivity list directors

Displays the inter-director connections for a director.

Contexts

`/:` connectivity list directors

`/engines/*/directors/*/:` list directors

Syntax

`connectivity list directors`

`[-h | --help]`

`[--verbose]`

`[-n | --directors] context path, context path...`

`[-f |] filename`

`[-u |] uuid`

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command
<code>[--verbose]</code>	Provides additional output during command execution. This may not have any effect for some commands.
<code>[-n --directors] <i>context path, context path ...</i></code>	Source director(s) for which connectivity should be reported.
<code>[-f --file] <i>filename</i></code>	Writes the connectivity report to the named file instead of echoing it. If the file exists, any previous contents will be lost. You can write the output to a file using an absolute path, or using a path relative to the CLI directory.
<code>[-d --uuid]</code>	Lists the connected directors by UUID instead of by name.

--	--

Description

Lists the other directors that are connected to the specified directors. The list includes the address, protocol, and local port name by which each remote director is connected to the specified directors.

See also

- `connectivity list all`
- `connectivity list initiators`
- `connectivity list storage-volumes`

connectivity list initiators

Displays the initiators connected to directors.

Contexts

`/:` `connectivity list initiators`
`/engines/*/directors/*/:` `list initiators`

Syntax

```
connectivity list directors
[-h | --help]
[--verbose]
[-n | --directors] context path, context path...
[-f | ] filename
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command
<code>[--verbose]</code>	Provides additional output during command execution. This may not have any effect for some commands.
<code>[-n --directors] <i>context path, context path ...</i></code>	Source director(s) for which connectivity should be reported.
<code>[-f --file] <i>filename</i></code>	Writes the connectivity report to the named file instead of echoing it. If the file exists, any previous contents will be lost. You can write the output to a file using an absolute path, or using a path relative to the CLI directory.
<code>[-d --uuid]</code>	Lists the connected directors by UUID instead of by name.

Description

Lists the initiators that are connected to a director. For each director specified, the list includes a table that reports each initiator's port WWN (FC initiators only) and node WWN (FC) or IQN (iSCSI), and to which port on the director they are connected.

See also

- `connectivity list all`
- `connectivity list directors`
- `connectivity list storage-volumes`

connectivity list storage-volumes

Displays the storage volumes connected to directors.

Contexts

`/:` `connectivity list storage-volumes`

`/engines/*/directors/*/:` `list storage-volumes`

Syntax

`connectivity list directors`

`[-h | --help]`

`[--verbose]`

`[-n | --directors] context path, context path...`

`[-f |] filename`

`[-s | --sort-by] key`

`[-l | --long-luns]`

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command
<code>[--verbose]</code>	Provides additional output during command execution. This may not have any effect for some commands.
<code>[-n --directors] <i>context path, context path ...</i></code>	Source director(s) for which connectivity should be reported.
<code>[-f --file] <i>filename</i></code>	Writes the connectivity report to the named file instead of echoing it. If the file exists, any previous contents will be lost. You can write the output to a file by using an absolute path, or by using a path relative to the CLI directory.
<code>[-s --sort-by] <i>key</i></code>	The field by which to sort the storage volume information (name, wwn or port).
<code>[-l --long-luns]</code>	Display LUNs as 16-digit hex-strings instead of as integers.

Description

Lists the storage volumes connected to a director. For each director, the list includes the address, protocol, and local port name by which each remote director is connected to the specified directors.

See also

- `connectivity list all`
- `connectivity list directors`
- `connectivity list initiators`

connectivity show

Displays the communication endpoints that can see each other.

Contexts

All contexts.

Syntax

```
connectivity show
[-p|--protocol[|ib|tcp|udp]
[e|--endpoints] port, port,...
```

Arguments

Optional arguments	
<code>[-p --protocol]</code> <code>{ib tcp udp}</code>	<p>Display endpoints with only the specified protocol. Arguments are case-sensitive, and include:</p> <ul style="list-style-type: none"> • <code>ib</code> - InfiniBand. Not supported in the current release. Use the <code>connectivity director</code> command to display IB protocol connectivity. • <code>tcp</code> - Transmission Control Protocol. • <code>udp</code> - UDP-based Data Transfer Protocol.
<code>[-e --endpoints]</code> <code>port,port...</code>	<p>List of one or more ports for which to display endpoints. Entries must be separated by commas. Default: Display endpoints for all ports.</p>

Description

Displays connectivity, but does not perform connectivity checks. Displays which ports can talk to each other.

See also

- `connectivity director`

connectivity validate-be

Checks that the back-end connectivity is correctly configured.

Contexts

All contexts.

Syntax

```
connectivity validate-be
[-d | --detailed]
```

[-h | --help]

--verbose

Arguments

Optional arguments	
[-h --help]	Displays the usage for this command.
[-d --detailed]	Details are displayed first, followed by the summary.
--verbose	Provides more output during command execution. This may not have any effect for some commands.

Description

This provides a summary analysis of the back-end connectivity information displayed by connectivity director if connectivity director was executed for every director in the system. It 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.
- Note:** 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.

Examples

Entering the `connectivity validate-be` command without any arguments provides a summary output as shown.

```
Vplexcli:/> connectivity validate-be
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*.
5019 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.
```

Display a summarized validation for back-end connectivity on an unhealthy system:

```

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*.
5019 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.

```

Display detailed validation for back-end connectivity on an unhealthy system:

```

VPlexcli:/> connectivity validate-be -d

Storage volumes that are not visible at all directors:

Cluster Director Storage Volumes Director Cannot See
-----
cluster-2 director-2-1-A VPD83T3:60001440000000103017bf5045090d75
VPD83T3:60001440000000103017bf5045090d70
VPD83T3:60001440000000103017bf5045092306
VPD83T3:60001440000000103017bf5045092f3b
VPD83T3:60001440000000103017bf5045092301
...
...
VPD83T3:6006016099751d002267f6538576e011
VPD83T3:60001440000000103017bf5045092310
VPD83T3:60001440000000103017bf5045090d6b
VPD83T3:60001440000000103017bf5045092f59
VPD83T3:60001440000000103017bf5045092f54
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*.
5019 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.

See also

- `connectivity director`
- `connectivity show`
- `connectivity validate-local-com`
- `connectivity validate-wan-com`
- `health-check`
- `validate-system-configuration`

connectivity validate-local-com

Validates that the actual connectivity over local-com matches the expected connectivity.

Contexts

All contexts.

Syntax

```
connectivity validate-local-com
[-c|--cluster] context path
[-e|--show-expected]
[-p|--protocol] communication protocol
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>context-path</i></code>	path of the cluster where local-com should be validated.
<code>[-e --show-expected]</code>	Prints the expected connectivity map instead of comparing it to the actual connectivity.
<code>[-p --protocol] <i>communication-protocol</i></code>	Specifies the protocol used for local-com (Fibre Channel or UDP). If not specified, the command attempts to determine the protocol based on the local-com ports in the system.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Verifies the expected local-com connectivity. This command assembles a list of expected local-com connectivity, compares it to the actual local-com connectivity, and reports any missing or extra connections. This command verifies only IP- or Fibre Channel-based local-com connectivity.

Expected connectivity is determined by collecting all ports whose role is local-com and verifying that each port in a port-group has connectivity to every other port in the same port-group.

When both Fibre Channel and IP ports with role local-com are present, the smaller subset is discarded and the protocol of the remaining ports is assumed to be the correct protocol.

See also

- `connectivity director`
- `connectivity show`
- `connectivity validate-be`
- `connectivity validate-wan-com`
- `health-check`
- `validate-system-configuration`

connectivity validate-wan-com

Verifies the expected IP and FC WAN COM connectivity.

Contexts

All contexts.

Syntax

```
connectivity validate-wan-com
  [-e|--show-expected]
  [-p|--protocol] communication-protocol
```

Arguments

Optional arguments	
<code>[-e --show-expected]</code>	Displays 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.
<code>[-p --protocol] <i>communication-protocol</i></code>	Specifies the protocol used for WAN COM (FC or UDP). If not specified, the command automatically selects either Fibre Channel or UDP based on the system’s Fibre Channel or Ethernet WAN COM ports.

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 Fibre Channel 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 port group at a cluster have connectivity to every other port in the same port group at all other clusters.

When both Fibre Channel 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.

See also

- `connectivity director`
- `connectivity show`
- `connectivity validate-be`

consistency-group add-virtual-volumes

Adds one or more virtual volume 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

Required arguments	
<code>[-v --virtual-volumes]</code> <i>virtual-volume, virtual-volume, ...</i>	* List of one or more comma-separated glob patterns or context paths of the virtual volume to add.
<code>[-g --consistency-group]</code> <i>consistency-group</i>	* Context path of the consistency group to which to add the specified virtual volume. If the current context is a consistency-group or below, then that consistency group is the default. Otherwise, this argument is required.

* - argument is positional.

Description

Adds the specified virtual volume to a consistency group. The properties of the consistency group immediately apply to the added volume.

Note: Only volumes with `visibility` and `storage-at-cluster` properties which match those of the consistency group can be added to the consistency group. Additionally, you cannot add a virtual volume to a consistency group if the initialization status of the virtual volume is `failed` or `in-progress`.

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.

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.

Note: 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.

Examples

Add multiple volumes using a single command. Separate virtual volume by commas:

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 volume 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
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                         synchronous
detach-rule                       active-cluster-wins
operational-status                 [(cluster-1,{ summary:: ok, details:: [] })), (cluster-2,{
summary:: ok, details:: [] })]
passive-clusters                   [cluster-1, cluster-2]
read-only                          false
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
- consistency-group list-eligible-virtual-volumes
- consistency-group remove-virtual-volumes
- *Dell EMC VPLEX Administration 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

Required arguments	
<code>[-c --cluster] <i>cluster</i></code>	*The cluster on which to roll back and resume I/O.
<code>[-g --consistency-group] <i>consistency-group</i></code>	* Context path of the consistency group on which to roll back and resume I/O.
Optional arguments	
<code>[-f --force]</code>	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 and resume I/O.

⚠ 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 are discarded.

Do not use this command to specify more than one cluster as the winner.

Examples

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` commands show a consistency group `my_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`.
 - Cluster-2, is still `suspended`, `cluster-departure`.
 - Cluster-1 is the winner, so it detached cluster-2.
- I/O at cluster-1 remains `suspended`, waiting for the administrator.

```

VPlexcli:/> ll /clusters/cluster-2/consistency-groups/
my_cg1/

/clusters/cluster-2/consistency-groups/my_cg1:

Attributes:
Name                               Value
-----
active-clusters                    []
cache-mode                         synchronous
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,
                                restore-link-or-choose-winner] })]

passive-clusters                  []
read-only                         false
recoverpoint-enabled              false
storage-at-clusters               []
virtual-volumes                   [dr1_read_write_latency_0000_12_vol]
visibility                        [cluster-1, cluster-2]

Contexts:
Name                               Description
-----
advanced                          -
recoverpoint                      -

VPlexcli:/clusters/cluster-2/consistency-groups/my_cg1> choose-winner -c
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

VPlexcli:/clusters/cluster-2/consistency-groups/my_cg1>
ls

Attributes:
Name                               Value
-----
active-clusters                    []
cache-mode                         synchronous
detach-rule                       no-automatic-winner
operational-status                [(cluster-1,{ summary:: suspended, details::
                                [cluster-departure, rebuilding-across-clusters] }),
                                (cluster-2,{ summary:: ok, details:: [] })]

passive-clusters                  []
read-only                         false
    
```

```

recoverpoint-enabled  false
storage-at-clusters  []
virtual-volumes      [dr1_read_write_latency_0000_12_vol]
visibility            [cluster-1, cluster-2]

```

```

Contexts:
advanced  recoverpoint

```

See also

- consistency-group resume-at-loser
- consistency-group summary
- *Dell EMC VPLEX Administration Guide*

consistency-group convert-to-local

Converts a distributed consistency group to a local consistency group.

context

All contexts

Syntax

```

convert-to-local
  [-h | --help]
  [--verbose]
  [[-c | --cluster=] cluster-context]
  [-f | --force]
  [[-g | --consistency-group=] consistency-group]

```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.
-c --cluster= <i>cluster context</i>	Specifies the cluster where all devices in the consistency group will be local.
-f --force	Forces the command to proceed, bypassing all user warnings .
-g --consistency-group= <i>consistency-group</i>	Specifies the consistency-group to make local.

Description

To convert a distributed consistency group to a local consistency group, this command converts all distributed devices under each virtual volume to local distributed devices. The legs on the specified cluster become the supporting device of the virtual volumes. All target devices should not be migration temporary devices and should not be exported to any other cluster.

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

Required arguments	
<code>[-n --name] <i>consistency-group name</i></code>	* Name of the new consistency group. Must be unique within a cluster. Name conflicts across s can be resolved by changing the name later using the <code>set name</code> command.
<code>[-c --cluster] <i>cluster</i></code>	Context path of the at which to create the consistency group. If the current context is a cluster or below, that 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 can be configured.

Each consistency group can contain up to 1000 .

All consistency groups have configurable properties that determine I/O behavior, including:

- `cache mode - synchronous`
- `visibility` - determines which s know about a consistency group. Default is only to the where the consistency group was created. Modified using the `set` command.
- `storage-at-clusters` - tells VPLEX at which 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 when there is an inter- 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- link is repaired after a failure. Default is true. Modified using the `set` command in `/clusters/cluster-n/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.
 - ① **Note:** 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.

Refer to the *Dell EMC Administration Guide* for more information about the consistency group properties.

Examples

In the following example,

- The `ls /clusters/*/consistency-groups/` command displays the names of all consistency groups in both s.
- 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.

- ① **Note:** See the Dell EMC 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_clwins 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_clwins 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
- consistency-group destroy
- consistency-group remove-virtual-volumes
- *Dell EMC VPLEX Administration Guide*

consistency-group destroy

Destroys the specified empty consistency groups.

Context

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	
<code>[-g --consistency-group] consistency-group, consistency-group, ...</code>	* List of one or more comma-separated context paths of the consistency groups to destroy.
Optional arguments	
<code>[-f --force]</code>	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 groups.

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.

Examples

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
- consistency-group remove-virtual-volumes
- *Dell EMC VPLEX Administration 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

Required arguments	
<code>[-g --consistency-group] <i>consistency-group</i></code>	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.
- Have the initialization status as `success`.

Examples

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
- consistency-group remove-virtual-volumes
- consistency-group summary
- *Dell EMC VPLEX Administration 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	
<code>[-v --virtual-volumes] <i>virtual-volume, virtual-volume, ...</i></code>	*Glob pattern or a list of one or more comma-separated context paths of the virtual volumes to remove from the consistency group.
<code>[-g --consistency-group] <i>context path</i></code>	*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.
<code>--force</code>	Do not ask for confirmation. 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.

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.

Note: 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.

Examples

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`
- `consistency-group destroy`
- *Dell EMC VPLEX Administration 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 synchronize 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                         synchronous
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                   []
read-only                          false
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           synchronous
detach-rule          no-automatic-winner
operational-status  [(cluster-1,{ summary:: ok, details:: [] } ),
                    (cluster-2,{ summary:: suspended, details:: [requires-
resume-at-loser] })]
passive-clusters     []
read-only            false
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-at-loser
- *Dell EMC VPLEX Administration 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	
<code>[-c --cluster] <i>cluster</i></code>	* The cluster on which to roll back and resume I/O.
<code>[-g --consistency-group] <i>consistency-group</i></code>	* The consistency group on which to resynchronize and resume I/O.
Optional arguments	
<code>[-f --force]</code>	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, you can 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, you 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.

You can then safely restart the applications at the losing cluster.

Without the `--force` option, this command asks for confirmation before proceeding, since its accidental use while applications are still running at the losing cluster could cause applications to misbehave.

Examples

```
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                           synchronous
detach-rule                          no-automatic-winner
operational-status                  [(cluster-1,{ summary:: ok, details:: [] } ),
                                (cluster-2,{ summary:: suspended, details:: [requires-
resume-at-loser] })]
passive-clusters                     []
read-only                            false
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                         synchronous
detach-rule                        no-automatic-winner
operational-status                 [(cluster-1,{ summary:: ok, details:: [] })),
                                (cluster-2,{ summary:: ok, details:: [] })]
passive-clusters                   []
read-only                          false
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
- consistency-group resume-after-rollback
- Dell EMC VPLEX Administration Guide

consistency-group set-detach-rule no-automatic-winner

Sets or changes 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,...
[-f|--force]
```

Arguments

Required arguments	
<code>[-g --consistency-group]</code> <i>consistency-group, consistency-group, ...</i>	The consistency groups on which to apply the no-automatic-winner detach rule.
Optional arguments	
<code>[-f --force]</code>	Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.

Description

Applies the no-automatic-winner detach rule to one or more specified consistency groups.

Note: This command requires user confirmation unless you use the `--force` argument.

This detach rule dictates no automatic detaches occur in the event of an inter-cluster link failure.

In the event of a cluster failure or departure, this rule-set results in I/O being suspended at all clusters whether or not VPLEX Witness is deployed. To resume I/O, use either the `consistency-group choose-winner` or `consistency-group resume-after-rollback` commands to designate the winning cluster.

Note: 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.

Examples

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 choose-winner`
- `consistency-group resume-after-rollback`
- `consistency-group set-detach-rule active-cluster-wins`
- `consistency-group set-detach-rule winner`
- Dell EMC VPLEX Administration Guide

consistency-group set-detach-rule winner

Sets the detach-rule for one or more synchronous consistency groups to winner. The specified cluster becomes the winner after the specified number of seconds.

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...
```

`[-f|--force]`

Arguments

Required arguments	
<code>[-c --cluster] <i>cluster-id</i></code>	The cluster that will be the winner in the event of an inter-cluster link failure.
<code>[-d --delay] <i>seconds</i></code>	The number of seconds after an inter-cluster link fails before the winning cluster detaches. Valid values for the delay timer are: <ul style="list-style-type: none"> 0 - Detach occurs immediately after the failure is detected. <i>number</i> - 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	
<code>[-g --consistency-group] <i>consistency-group, consistency-group, ...</i></code>	The consistency groups on which to apply the winner detach rule.
<code>[-f --force]</code>	Force the operation to continue without confirmation. Allows this command to be run from non-interactive scripts.

Description

Applies the winner detach rule to one or more specified synchronous consistency groups.

Note: This command requires user confirmation unless the `--force` argument is used.

In the event of a cluster failure or departure, this rule-set results in I/O continuing on the selected cluster only. I/O will be suspended at all other clusters. If VPLEX Witness is deployed it will override this selection if the selected cluster has failed.

Note: 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.

Examples

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
- consistency-group set-detach-rule no-automatic-winner
- *Dell EMC VPLEX Administration 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      Operational Status      Status      Active      Passive
Name                  -----
-----
-----
GRP3                   Suspended
0 consistency groups have risk of roll back.
```

See also

- consistency-group create
- consistency-group destroy
- *Dell EMC VPLEX Administration Guide*

date

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

Contexts

All contexts.

Syntax

```
date
```

Examples

```
VPlexcli:/> date
Tue Jul 20 15:57:55 UTC 2010director ping
```

describe

Describes the attributes of the given context.

Contexts

All contexts with attributes.

Syntax

describe

`[-c|--context] context-path`

Arguments

Optional arguments	
<code>[-c --context] context-path</code>	Context to describe.

Examples

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|device-name}
[-m|--mirror] {context-path|mirror-name}
[-r|--rule-set] rule-set
[-f|--force]
```

Arguments

Required arguments	
<code>[-d --device]</code> <i>context-path</i> or <i>device-name</i>	* 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.
Optional arguments	
<code>[-m --mirror]</code> <i>context-path</i> or <i>mirror-name</i>	* Name or context path of the mirror to detach. Does not need to be a top-level device. If the name of a device is used, ensure the device name is not ambiguous. For example, ensure that the same device name is not used by local devices in different clusters.
<code>[-r --rule-set]</code> <i>rule-set</i>	Rule-set to apply to the distributed device that is created when a mirror is added to a local device.
<code>[-f --force]</code>	When <code>--force</code> is set, do not ask for confirmation when attaching a mirror. Allows this command to be run using a non-interactive script. If the <code>--force</code> 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.

* - argument is positional.

Description

If the parent device is a RAID 0 or RAID C, it is converted to a 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 you do not specify the `--rule-set` argument, 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.

When attaching a remote mirror to a local device, or when attaching a new mirror to a distributed device, both operations consume slots. Both scenarios result in the same following error message:

```
VPlexcli:/clusters/cluster-2/storage-elements/extents> device attach-mirror -
d rhyan_dr_test -m rhyan_mig_src_0000
device attach-mirror: Evaluation of <<device attach-mirror -d dr_test -m
mig_src_0000>> failed.
cause: Unable to attach mirror 'mig_src_0000' to device
'dr_test'.
cause: Unable to attach mirror 'mig_src_0000' to distributed
Device 'dr_test'.
cause: Firmware command error.
cause: Active metadata device does not have a free slot.
```

Refer to the troubleshooting section of the VPLEX procedures in the SolVe Desktop for instructions on increasing the number of slots.

i **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.

- The rule set that will be applied to the new distributed device potentially allows conflicting detaches.

Homogeneous array requirement for thin volumes

To preserve thinness of the new RAID-1 device where the parent device is created on a thin volume and is thin-capable, the mirror device must be created from the same storage-array-family as the parent device. If the user tries to attach a mirror leg from a dissimilar array-family, the command displays a warning that the thin-capability of the RAID-1 device will be lost and it can render the virtual volume to be thin disabled. The following is an example of the warning message:

```
VPlexcli:/> device attach-mirror --device xio_device --mirror vnx_device
Thin-capability is only supported with homogeneous storage-array types. The
top-level device
'xio_device' is supported by XtremIO but the mirror 'vnx_device' is supported
by CLARiiON.
Since XtremIO and CLARiiON are not homogeneous, the top-level device will
lose thin-capability
after the new mirror is attached. Do you wish to proceed? (Yes/No) No

device attach-mirror: Evaluation of <<device attach-mirror --device
xio_device --mirror vnx_device>>
failed.
cause: Unable to attach mirror 'vnx_device' to device
'xio_device'.
cause: Operation was halted by the user

VPlexcli:/>
```


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
- device detach-mirror
- rebuild status

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|device-name]
```

Arguments

Required arguments	
[-d --device] [<i>context-path</i> <i>device-name</i>]	* 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:



Applicable to one-legged devices that are not top-level.

Examples

```

VPlexcli:/clusters/cluster-1/devices> device collapse --device /clusters/
cluster-1/devices/A
  
```

device detach-mirror

Removes (detaches) a mirror from a RAID-1 device.

Contexts

All contexts.

Syntax

```

device detach-mirror
[-d|--device] [context-path|device-name]
[-m|--mirror] [context-path|mirror-name]
[-s|--slot] slot-number
[-i|--discard]
[-f|--force]
  
```

Arguments


Required arguments	
<code>[-d --device]</code> <i>context-path or device-name</i>	* 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 system, including local devices on other clusters.
Optional arguments	
<code>[-m --mirror]</code> <i>context-name or mirror-name</i>	* 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.
<code>[-s --slot]</code> <i>slot-number</i>	Slot number of the mirror to be discarded. Applicable only when the <code>--discard</code> argument is used.
<code>[-i --discard]</code>	When specified, discards the detached mirror. The data is not discarded.
<code>[-f --force]</code>	Force the mirror to be discarded. Must be used when <code>--discard</code> argument is used. The <code>--force</code> argument is set for detaching an unhealthy mirror it is discarded even if <code>--discard</code> flag is not set.

* - argument is positional.

Description

Use this command to detach a mirror leg from a RAID 1 device.


Figure 1 RAID device and virtual volume: before detach mirror



If the RAID device supports a virtual volume, and you don't use the `--discard` argument the command:

- Removes the mirror (child device) from the RAID 1 parent device.
- Makes the detached child a top-level device.
- Creates a new virtual volume on top the new device and prefixes the name of the new device with the name of the original device.

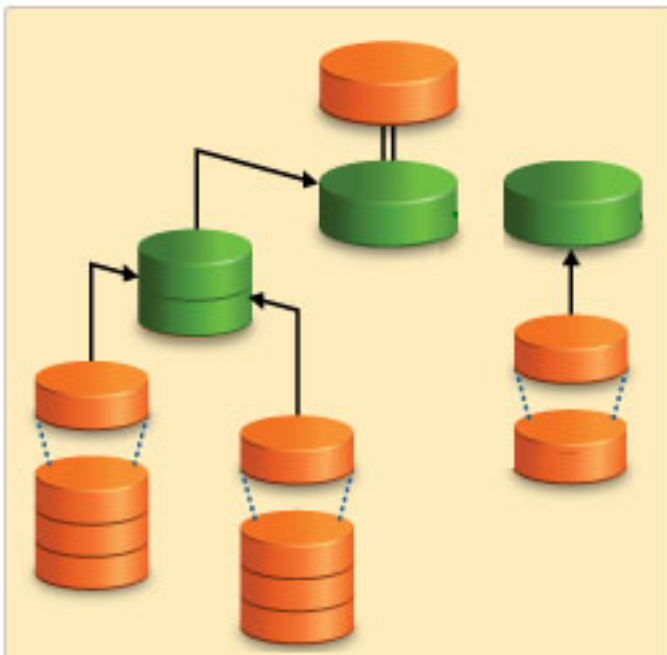
Figure 2 Devices and virtual volumes: after detach mirror - no discard



If the RAID device supports a virtual volume, and you use the `--discard` argument, the command:

- Removes the mirror (child device) from the RAID 1 parent device.
- Makes the detached child a top-level device.
- Creates no new virtual.
- Detaches the mirror regardless of its current state and does not guarantee data consistency.

Figure 3 Devices and virtual volumes: after detach mirror - with discard



See also

- `device attach-mirror`

device mirror-isolation auto-unisolation disable

Disables mirror auto-unisolation.

Contexts

All contexts.

Syntax

```
device mirror-isolation auto-unisolation disable
[-f|--force]
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-f --force]</code>	Forces the operation to continue without confirmation.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Mirror isolation provides a mechanism to automatically unisolate mirrors that were previously isolated. When mirror isolation feature is enabled, disabling mirror auto-unisolation will prevent the system from automatically unisolating any isolated mirrors whose underlying storage-volume's performance is now in the acceptable range.

For the option to manually unisolate the mirror, follow the troubleshooting procedure for VPLEX in the SolVe Desktop.

Examples

Shows the result when the command is executed on all clusters when mirror isolation is disabled:

```
Vplexcli:/> device mirror-isolation auto-unisolation disable
Mirror isolation provides a mechanism to automatically unisolate mirrors that
were previously isolated. This operation will prevent the system from
automatically unisolating the underlying storage-volumes once their
performance is in the acceptable range. You can manually unisolate the mirror
by following the troubleshooting procedure.
Continue to disable auto-unisolation? (Yes/No) y
Auto-unisolation is disabled on clusters cluster-1,cluster-2.
```

Shows the command executed with the `--force` option, when the mirror-isolation feature is disabled:

```
Vplexcli:/> device mirror-isolation auto-unisolation disable -f
Mirror isolation is not enabled on clusters cluster-1,cluster-2.
Mirror isolation provides a mechanism to automatically unisolate mirrors that
were previously isolated. When mirror isolation is enabled, this operation
```

will prevent the system from automatically unisolating the underlying storage volumes once their performance is in the acceptable range. You can manually unisolate the mirror by following the troubleshooting procedure. Auto-unisolation is disabled on clusters cluster-1,cluster-2.

Shows auto-unisolation was not disabled because the feature is not supported:

```
VPlexcli:/> device mirror-isolation auto-unisolation disable
Mirror isolation is not enabled on clusters cluster-1,cluster-2.
Mirror isolation provides a mechanism to automatically unisolate mirrors that
were previously isolated. When mirror isolation is enabled, this operation
will prevent the system from automatically unisolating the underlying storage-
volumes once their performance is in the acceptable range. You can manually
unisolate the mirror by following the troubleshooting procedure.
Continue to disable auto-unisolation? (Yes/No) y
device mirror-isolation auto-unisolation disable: Evaluation of <<device
mirror-isolation
auto-unisolation disable>> failed.
cause: Could not disable auto unisolation.
cause: Could not disable auto unisolation.
Modifying auto unisolation is not
supported by this version of firmware.
```

Shows auto-unisolation failed because one cluster is not available:

```
VPlexcli:/> device mirror-isolation auto-unisolation disable
device mirror-isolation auto-unisolation disable: Evaluation of <<device
mirror-isolation auto-unisolation disable>> failed.
cause: Could not disable auto unisolation.
cause: Firmware command error.
cause: communication error recently.
```

See also

- device mirror-isolation auto-unisolation enable
- device mirror-isolation disable
- device mirror-isolation enable
- device mirror-isolation show
- *Dell EMC VPLEX Administration Guide*
- Dell EMC VPLEX Procedures in SolVe Desktop

device mirror-isolation auto-unisolation enable

Enables mirror auto-unisolation.

Contexts

All contexts.

Syntax

```
device mirror-isolation auto-unisolation enable
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

This command enables auto mirror unisolation.

Mirror isolation provides a mechanism to automatically unisolate mirrors that were previously isolated. When mirror isolation is enabled, auto-unisolation allows the system to automatically unisolate the underlying storage-volumes once their performance is in the acceptable range.

Examples

Shows auto-unisolation enabled when mirror isolation is disabled on both clusters:

```
VPlexcli:/> device mirror-isolation auto-unisolation enable
Mirror isolation is not enabled on clusters cluster-1,cluster-2.
Auto-unisolation is enabled on clusters cluster-1,cluster-2.
```

Shows auto-unisolation enabled when mirror isolation is disabled on one of the clusters:

```
VPlexcli:/> device mirror-isolation auto-unisolation enable
Mirror isolation is not enabled on cluster cluster-2.
Auto-unisolation is enabled on clusters cluster-1,cluster-2.
```

Shows auto-unisolation when mirror isolation is enabled:

```
VPlexcli:/> device mirror-isolation auto-unisolation enable
Auto-unisolation is enabled on clusters cluster-1,cluster-2.
```

Shows auto-unisolation enable operation failed as the feature is not supported:

```
VPlexcli:/> device mirror-isolation auto-unisolation enable
Mirror isolation is not enabled on clusters cluster-1,cluster-2.
device mirror-isolation auto-unisolation enable: Evaluation of <<device
mirror-isolation auto-unisolation enable>> failed.
cause:                               Could not enable auto unisolation.
cause:                               Could not enable auto unisolation.
                                     Modifying auto unisolation is not supported by
this version of firmware.
```

Shows auto-unisolation enable operation failed because one cluster is not available:

```
VPlexcli:/> device mirror-isolation auto-unisolation enable
device mirror-isolation auto-unisolation enable: Evaluation of <<device
mirror-isolation auto-unisolation enable>> failed.
cause:                               Could not enable auto
unisolation.
cause:                               Firmware command error.
cause:                               communication error
recently.
```


Shows auto-unisolation enable operation failed because the meta volume is not ready:

```

VPlexcli:/> device mirror-isolation auto-unisolation enable
device mirror-isolation auto-unisolation enable: Evaluation of <<device
mirror-isolation auto-unisolation enable>> failed.
cause: Could not enable auto
unisolation.
cause: Could not enable auto
unisolation.
Firmware command error.
The active metadata device is not healthy enough to persist the change fully.
    
```

See also

- device mirror-isolation auto-unisolation disable
- device mirror-isolation disable
- device mirror-isolation enable
- device mirror-isolation show
- Dell EMC VPLEX Administration Guide

device mirror-isolation disable

Disables mirror isolation on the specified clusters.

Contexts

All contexts.

Syntax

```

device mirror-isolation disable
[-c|--clusters] context-path [, context-path...]
[-f|--force]
[-h|--help]
[--verbose]
    
```

Arguments

Optional arguments	
<code>[-c --clusters] <i>context-path</i> [, <i>context-path...</i>]</code>	Specifies the list of clusters on which to disable mirror isolation.
<code>[-f --force]</code>	Forces the operation to continue without confirmation.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

A RAID 1 mirror leg built upon a poorly performing storage volume can bring down the performance of the whole RAID 1 device and increase I/O latencies to the applications using this device. VPLEX prevents I/Os to these poorly performing mirror legs to improve the RAID 1 performance. This behavior or feature is known as mirror isolation.

When disabling the mirror isolation feature on one or more clusters, this command prints a warning and asks for confirmation.

Note: This command disables the mirror isolation feature and prevents VPLEX from improving the performance of a RAID 1 device containing a poorly performing mirror leg. This command should only be used if redundancy is desired over RAID 1 performance improvement.

Examples

Disable mirror isolation on all clusters:

```
Vplexcli:/> device mirror-isolation disable
Disabling the mirror isolation feature will prevent VPLEX from improving the
performance of a RAID-1 device containing a poorly performing mirror leg.
This command should be only used if redundancy is desired over RAID-1
performance improvement.
Continue to disable mirror isolation on cluster-1,cluster-2? (Yes/No) Yes
Mirror isolation has been disabled on 'cluster-1'.
Mirror isolation has been disabled on 'cluster-2'.
```

Disable mirror isolation on all clusters without being prompted to confirm:

```
Vplexcli:/> device mirror-isolation disable -f
WARNING: Disabling the mirror isolation feature will prevent VPLEX from
improving the performance of a RAID-1 device containing a poorly performing
mirror leg. This command should be only used if redundancy is desired over
RAID-1 performance improvement.
Mirror isolation has been disabled on 'cluster-1'.
Mirror isolation has been disabled on 'cluster-2'.
```

Disable mirror isolation on one cluster:

```
Vplexcli:/> device mirror-isolation disable -c cluster-1
Disabling the mirror isolation feature will prevent VPLEX from improving the
performance of a RAID-1 device containing a poorly performing mirror leg.
This command should be only used if redundancy is desired over RAID-1
performance improvement.
Continue to disable mirror isolation on cluster-1 (Yes/No) Yes
Mirror isolation has been disabled on 'cluster-1'.
```

Attempt to disable mirror-isolation on the clusters when it is already disabled:

```
Vplexcli:/> device mirror-isolation disable -f
WARNING: Disabling the mirror isolation feature will prevent VPLEX from
improving the performance of a RAID-1 device containing a poorly performing
mirror leg. This command should be only used if redundancy is desired over
RAID-1 performance improvement.
Mirror isolation has been disabled on 'cluster-1'.
Mirror isolation has been disabled on 'cluster-2'.
```

Attempt to disable mirror-isolation on a system where mirror-isolation is not supported:

```
Vplexcli:/> device mirror-isolation disable -f
WARNING: Disabling the mirror isolation feature will prevent VPLEX from
improving the performance of a RAID-1 device containing a poorly performing
mirror
leg. This command should be only used if redundancy is desired over
RAID-1 performance improvement.
device mirror-isolation disable: Evaluation of <<device mirror-isolation
```

```

disable -f>> failed.
cause:          Could not disable mirror isolation on
'cluster-1'.
cause:          Firmware command error.
cause:          no such command.

```

Attempt to disable mirror-isolation on both clusters and succeeded on cluster 1, but failed on cluster 2 because the feature is not supported:

```

VPlexcli:/> device mirror-isolation disable
Disabling the mirror isolation feature will prevent VPLEX from improving the
performance of a RAID-1 device containing a poorly performing mirror leg.
This command should be only used if redundancy is desired over RAID-1
performance improvement.
Continue to disable mirror isolation on cluster-1,cluster-2 (Yes/No) Yes
Mirror isolation has been disabled on 'cluster-1'.
device mirror-isolation disable: Evaluation of <<device mirror-isolation
disable>> failed.
cause:          Could not disable mirror isolation on
'cluster-2'.
cause:          Firmware command error.
cause:          invalid subcommand.

```

See also

- `device mirror-isolation auto-unisolation disable`
- `device mirror-isolation auto-unisolation enable`
- `device mirror-isolation enable`
- `device mirror-isolation show`
- *Dell EMC VPLEX Administration Guide*

device mirror-isolation enable

Enables mirror isolation on the specified clusters.

Contexts

All contexts.

Syntax

```

device mirror-isolation enable
[-c|--clusters] context-path [, context-path...]
[--also-enable-autounisolation]
[-h|--help]
[--verbose]

```


Arguments

Optional arguments	
<code>[-c --clusters] <i>context-path</i> [, <i>context-path</i>...]</code>	Specifies the list of clusters on which to enable mirror isolation.
<code>[--also-enable-autounisolation]</code>	Enables auto-unisolation if specified.

<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

A RAID 1 mirror leg built on a poorly performing storage volume can bring down the performance of the whole RAID 1 device and increase I/O latencies to the applications using this device. VPLEX prevents I/Os to these poorly performing mirror legs to improve the RAID 1 performance. This behavior or feature is known as mirror isolation.

 **Note:** This command enables the mirror isolation feature and should only be used if RAID 1 performance improvement is desired over redundancy.

Examples

Enable mirror isolation on all clusters:

```
VPlexcli:/> device mirror-isolation enable
Mirror isolation has been enabled on 'cluster-1'.
Mirror isolation has been enabled on 'cluster-2'.
```

Or

```
VPlexcli:/> device mirror-isolation enable -c *
Mirror isolation has been enabled on 'cluster-1'.
Mirror isolation has been enabled on 'cluster-2'.
```

Enable mirror isolation on one cluster:

```
VPlexcli:/> device mirror-isolation enable -c cluster-1
Mirror isolation has been enabled on 'cluster-1'.
```

Attempt to enable mirror-isolation on the clusters when it is already enabled:

```
VPlexcli:/> device mirror-isolation enable
Mirror isolation has been enabled on 'cluster-1'.
Mirror isolation has been enabled on 'cluster-2'.
```

Enable mirror-isolation when auto-unisolation is disabled:

```
VPlexcli:/> device mirror-isolation enable
Mirror isolation has been enabled on 'cluster-1'.
Mirror isolation has been enabled on 'cluster-2'.
Please be aware that auto unisolation is disabled. In order to manually
enable this feature you can use 'device mirror-isolation auto-unisolation
enable'.
```

Enable mirror-isolation and auto-unisolation on both clusters:

```
VPlexcli:/> device mirror-isolation enable --also-enable-autounisolation
Mirror isolation has been enabled on 'cluster-1'.
```

```
Mirror isolation has been enabled on 'cluster-2'.
Auto-unisolation is enabled on clusters cluster-1,cluster-2.
```

Attempt to enable mirror-isolation on a system where mirror-isolation is not supported:

```
VPlexcli:/> device mirror-isolation enable
device mirror-isolation enable: Evaluation of <<device mirror-isolation
enable>> failed.
cause:                               Could not enable mirror isolation on
'cluster-1'.
cause:                               Firmware command error.
cause:                               no such command.
```

Attempt to enable mirror-isolation on both clusters and succeeded on cluster 1, but failed on cluster 2 because the feature is not supported:

```
VPlexcli:/> device mirror-isolation enable
Mirror isolation has been enabled on 'cluster-1'.
device mirror-isolation disable: Evaluation of <<device mirror-isolation
enable>> failed.
cause:                               Could not enable mirror isolation on
'cluster-2'.
cause:                               Firmware command error.
cause:                               invalid subcommand.
```

See also

- `device mirror-isolation auto-unisolation disable`
- `device mirror-isolation auto-unisolation enable`
- `device mirror-isolation disable`
- `device mirror-isolation show`
- Dell EMC VPLEX Administration Guide

device mirror-isolation show

Lists the configuration parameters related to mirror isolation for the specified clusters.

Contexts

All contexts.

Syntax

```
device mirror-isolation show
[-c|--clusters] context-path [, context-path...]
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --clusters] <i>context-path</i> [, <i>context-path...</i>]</code>	Specifies the list of clusters on which to show mirror isolation configuration parameters.

[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

Used to display all the configuration parameters related to mirror isolation for the specified clusters.

The current configuration parameters supported are:

Table 13 Supported configuration parameters

Enabled	Indicates "true" if the feature is enabled, "false" if disabled and "<not available>" if the value could not be retrieved.
Auto Unisolation	Indicates "true" if the system will automatically unisolate an isolated mirror when the underlying storage-volume's performance is in the acceptable range, "false" if manual unisolation was desired, and "<not available>" if the value could not be retrieved.
Isolation Interval	Indicates the isolation sweep interval in seconds if the value was retrieved successfully, and "<not available>" if the value could not be retrieved.
Unisolation interval	Indicates the unisolation sweep interval in seconds if the value was retrieved successfully, and "<not available>" if the value could not be retrieved.

If a value for any configuration parameter cannot be retrieved for the cluster, it may be because the feature is not supported or there was a command failure.

Examples

Shows the mirror isolation configuration parameters on all clusters:

```

Vplexcli:/> device mirror-isolation show
Cluster      Enabled  Auto unisolation  Isolation Interval  Unisolation
Interval
-----
cluster-1   true    false             60                   14400
cluster-2   true    false             60                   14400
    
```

Or

```

Vplexcli:/> device mirror-isolation show -c *
Cluster      Enabled  Auto unisolation  Isolation Interval  Unisolation
Interval
-----
cluster-1   true    false             60                   14400
cluster-2   true    false             60                   14400
    
```

Shows the mirror isolation configuration parameters on one cluster:

```
VPlexcli:/> device mirror-isolation show -c cluster-1
Cluster      Enabled  Auto unisolation  Isolation Interval  Unisolation
Interval
-----
cluster-1   true    false             60                  14400
```

Shows the command executed on an invalid cluster:

```
VPlexcli:/> device mirror-isolation show -c blah
device mirror-isolation show:      Evaluation of <<device mirror-isolation
show -c blah>> failed.
cause:                             Command execution failed.
cause:                             Failed to set value for option --clusters.
cause:                             Could not find appropriate contexts matching
'[blah]'.
```

Shows the firmware command or feature is not supported:

```
VPlexcli:/> device mirror-isolation show
Cluster      Enabled  Auto unisolation  Isolation Interval  Unisolation
Interval
-----
cluster-1   <not available> <not available> <not available> <not
available>
cluster-2   <not available> <not available> <not available> <not
available>
```

Shows the auto-unisolation feature is not supported:

```
VPlexcli:/> device mirror-isolation show
Cluster      Enabled  Auto unisolation  Isolation Interval  Unisolation
Interval
-----
cluster-1   true    <not available>   60                  14400
cluster-2   true    <not available>   60                  14400
```

See also

- `device mirror-isolation auto-unisolation disable`
- `device mirror-isolation auto-unisolation enable`
- `device mirror-isolation disable`
- `device mirror-isolation enable`
- *Dell EMC VPLEX Administration Guide*

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
[-a|--all-at-island]
[-f|--force]
```

Arguments

Optional arguments	
<code>[-c --cluster]</code> <i>context path</i>	Resume I/O on the specified cluster and the clusters it is in communication with during a link outage. Applicable only when the <code>all-at-island</code> argument is used or when the specified devices are distributed devices. Not required for local devices with global visibility.
<code>[-r --devices]</code> <i>context path or device-name</i>	Name or context path of the devices for which to resume I/O. They must be top-level devices.
<code>[-a --all-at-island]</code>	Resume I/O on all devices on the chosen winning cluster and the clusters with which it is communicating.
<code>[-f --force]</code>	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 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 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.

Examples

```
VPlexcli:/distributed-storage/distributed-devices> device resume-link-down --
all-at-island --cluster --devices DD_5d --force
```

See also

- `device resume-link-up`
- `ds dd declare-winner`

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

Optional arguments	
<code>[-r --devices] <i>context path, context path...</i></code>	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.
<code>[-v --virtual-volume] <i>context path,context path...</i></code>	Resume I/O on the specified virtual volumes.
<code>[-a --all]</code>	Resume I/O on all virtual volumes on the losing cluster.
<code>[-f --force]</code>	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 mirror-isolation disable`

device resurrect-dead-storage-volumes

Resurrect the thin-aware storage-volumes supporting the target devices that are marked dead.

Contexts

Any

Syntax

```
device resurrect-dead-storage-volumes
[-h|--help]
[--verbose]
[-r|--devices=device [,device [,device]]
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.
Required arguments	
-r --devices= <i>device</i> [, <i>device</i> [, <i>device</i>]]	*Specifies the devices to resurrect dead supporting storage-volumes on. The <i>device</i> name can include wild card symbols.

* Argument is positional.

Description

This command is used for storage volumes that do not auto-resurrect after they receive an Out Of Space error on a write command and become hardware dead. This scenario should only happen on an XtremIO storage volume that is used as a VPLEX mirror leg. After resolving the underlying issue that lead to an out of space error, use this command to resume I/O for supporting storage-volumes that have been marked dead. The target devices may be of any geometry, local or distributed. This command executes `storage-volume resurrect` for all dead-storage-volumes of a device. This scenario should only happen on an XtremIO storage volume that is used as a VPLEX mirror leg.

director appcon

Runs the application console on Linux systems.

Contexts

All

Syntax

```
director appcon
--xterm-opts options
[-t|--targets] target-glob,target-glob...
--timeout seconds
--show-plan-only
```

Arguments

Optional arguments	
<code>--xterm-opts</code>	Additional options for xterm on Linux. Default: <code>-sb -sl 2000</code> .
<code>[-t --targets]</code> <i>target-glob, target-glob...</i>	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.
<code>--timeout</code> <i>seconds</i>	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.
<code>--show-plan-only</code>	Shows the targets that will be affected, but the actual operation is not performed. Recommended when the <code>--targets</code> argument is used.

Description

Applicable only to Linux systems.

Opens the hardware application console for each director in a separate window.

Examples

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`

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

Optional arguments	
<code>[-d --dir] <i>directory</i></code>	Directory in which to put dump files. Default: Current directory.
<code>[-s --no-timestamp]</code>	Omits the timestamp in the dump filenames.
<code>[-o --overwrite]</code>	Overwrites existing files. Default: Not to overwrite.
<code>[-c --include-cores]</code>	Includes the director's core image in the dump.


	 WARNING Including core images can make this command take a very long time.
<code>[-z --no-zip]</code>	Turns off the packaging of dump files into a compressed zip file.
<code>[-p --no-progress]</code>	Turns off progress reporting.
<code>[-t --targets]</code> <i>target-glob, target-glob...</i>	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.
<code>--timeout <i>seconds</i></code>	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.
<code>--show-plan-only</code>	Shows the targets that will be affected, but the actual operation is not performed. Recommended when the <code>--targets</code> argument is used.

Description

Used by automated scripts and by Dell 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.

Examples

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
- collect-diagnostics
- getsysinfo
- smsdump

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

Optional arguments	
<code>[-t --targets]</code> <i>target-glob, target-glob...</i>	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.
<code>--timeout</code> <i>seconds</i>	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.
<code>--show-plan-only</code>	Shows the targets that will be affected, but the actual operation is not performed. Recommended when the <code>--targets</code> argument is used.

Description

Used by automated scripts and by Dell EMC Customer Support to help troubleshoot problems.

Examples

```
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`

- `director appdump`

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
[-t|--timeout] seconds
[-a|--apply-cluster-settings]
[-f|--force]
```

Arguments

Required arguments	
<code>[-n --director] <i>director</i></code>	* The director to be commissioned.
Optional arguments	
<code>[-f --force]</code>	Commission the director regardless of firmware version mismatch.
<code>--timeout <i>seconds</i></code>	The maximum time to wait for <code>--apply-cluster-settings</code> operations to complete, in seconds. Default: 60 seconds. 0: No timeout.
<code>[a --apply-cluster-settings]</code>	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.

Examples

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`
- `version`

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	
<code>[-n --director] <i>director</i></code>	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.

Examples

```
VPlexcli: /> director decommission --director Cluster_1_Dir1A
```

See also

- `director commission`
- `director forget`
- `director shutdown`

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 context, command is `fc-port-stats`

Syntax

```
director fc-port-stats
```

```
[-d|--director] director
```


`[-o|--role] role`

`[-r|--reset]`

Arguments

Required arguments	
<code>[-d --director] <i>director</i></code>	Context path of the director for which to display FC statistics. Not required if the current context is <code>/engines/engine/directors/director</code> .
Optional arguments	
<code>[-o --role] <i>role</i></code>	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 <code>--reset</code> is specified. Roles include: <ul style="list-style-type: none"> <code>back-end</code> - Filter on ports used to access storage devices that the system itself does I/O to. <code>front-end</code> - Filter on ports used to make storage available to hosts. <code>inter-director-communication</code> - Filter on ports used to communicate with other directors. <code>local-com</code> - Filter on ports used to communicate with other directors at the same cluster. <code>management</code> - Filter on ports used to communicate with the management server. <code>wan-com</code> - Filter on ports used to communicate with other clusters.
<code>[-r --reset]</code>	Reset the statistics counters of all ports at the specified director. If you specify this argument, the command ignores the <code>--role</code> argument.

Description

Displays statistics generated by the driver for FibreChannel ports at the specified director and optionally with the specified role, or resets those statistics.

Run this command from the `/engines/engine/directors/director` context to display the Fibre Channel statistics for the director in the current context.

Examples

Display a director's Fibre Channel port statistics from the root context:

```
VPlexcli:/> director fc-port-stats -d director-2-1-A
```

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
```

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
0
- Expired:     0          0          0          0          0          0
- Bad CRCs:    0          0          0          0          0          0
- Encoding Errors: 0          0          0          0          0          0
- Out Of Order: 0          0          0          0          0          0
0
- Lost:        0          0          0          13         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
    
```

See also

- `monitor stat-list`

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

Optional arguments	
<code>[-t --targets]</code> <i>director,director...</i>	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.

Table 14 director firmware show-banks field descriptions

Field	Description
Banks	Each director has two firmware banks; A and B.
Status	active - The software in this bank is currently operating on the director. inactive - The software in this bank is not operating on the director.
Marked for next reboot	no - The software in this bank will not be used the next time the director reboots. yes - The software in this bank will be used the next time the director reboots.
Director Software version	Software version currently operating in the director.

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

```

See also

- version

director forget

Removes a director from the VPLEX.

Contexts

All contexts.

Syntax

```
director forget
```

```
[-n|--director] director uuid
```

Arguments

Required arguments

<code>[-n --director] <i>director-uuid</i></code>	Director ID number. Use the <code>ll</code> command in <code>engines/engine/directors</code> context to display director ID numbers.
---	--

Description

Removes the specified director from the context tree. Deletes all information associated with the director.

Examples

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  -----             ID           -----      Status
Status
-----             -----             -----      -----
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`
- `director decommission`

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	
--------------------	--

<code>[-n --director] <i>director</i></code>	The remote director on which to change the access password.
<code>[-c --current-password] <i>current-password</i></code>	The current access password of the specified director.
<code>[-p --new-password] <i>new-password</i></code>	The new access password to set for the specified director.

Description

Changes the password for a specified director.

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	
<code>[-i --ip-address] <i>IP-address</i></code>	The target's IP address.
Optional arguments	
<code>[-n --director] <i>director</i></code>	The director from which to perform the operation.
<code>[-w --wait] <i>seconds</i></code>	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.

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.

Examples

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`

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	
<code>[-f --force]</code>	Forces this operation.
Optional arguments	
<code>[-n --director] <i>context-path</i></code>	* 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.

Examples

In the following example:

- The `director shutdown` command shuts down DirA.
- The `ll` command displays the shutdown director.

```

VPlexcli:/engines/engine-1-1/directors/director-1-1-A> director shutdown
--force
Please check the status of the director: director-1-1-A for its shutdown
status.
Status      Description
-----
Started.    Shutdown started.
VPlexcli:/engines/engine-1-1/directors/director-1-1-A> ll
Attributes:
Name                                     Value
-----
auto-boot                                 true
auto-restart                              true
.
.
.
marker-led                                off
operational-status                        stopped
.
.
.

```

See also

- `cluster shutdown`
- `director commission`

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	
<code>[-i --ip-address] <i>IP-address</i></code>	The target's IP address. This address is one of the Ethernet WAN ports on another director. Use the <code>ll port-group</code> command to display the Ethernet WAN ports on all directors.
Optional arguments	
<code>[-n --director] <i>director</i></code>	The name of the director from which to perform the operation. Can be either the director's name (for example <code>director-1-1-A</code>) 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 fire wall 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.

 **CAUTION** If the target machine does not respond properly, the `traceroute` might stall. Run this command multiple times.

See also

- `director ping`

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.

Examples

Shows director uptime:

```
VPlexcli:/> director uptime
Director director-1-2-A: 0 days, 0 hours, 52 minutes, 29 seconds.
Director director-1-2-B: 0 days, 0 hours, 52 minutes, 11 seconds.
Director director-2-2-A: 0 days, 0 hours, 52 minutes, 19 seconds.
Director director-2-1-A: 0 days, 0 hours, 52 minutes, 34 seconds.
Director director-2-2-B: 0 days, 0 hours, 52 minutes, 31 seconds.
Director director-1-1-A: 0 days, 0 hours, 57 minutes, 2 seconds.
```



```
Director director-1-1-B: 0 days, 0 hours, 52 minutes, 24 seconds.
Director director-2-1-B: 0 days, 0 hours, 52 minutes, 28 seconds.
```

See also

- cluster shutdown
- director firmware show-banks

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.

Examples

```
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

disconnect

Disconnects one or more connected directors.

Contexts

All contexts.

Syntax

```
disconnect
```


```
[-n|--directors] context-path, context-path...
```

Arguments

Required arguments	
[-n --directors] <i>context-path, context-path...</i>	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 directors.

This command is used in various procedures in the *EMC VPLEX Troubleshooting Guide*.

Examples

```
Vplexcli:> disconnect --directors /engines/engine-1-1/directors/director-1-1-B
```

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
[-m|--migrations] context-path,context-path...
[-f|--force]
```

Arguments

Required arguments	
<code>[-m --migrations]</code>	* List of one or more migrations to cancel. Entries must be separated by commas.
Optional arguments	
<code>[-f --force]</code>	Forces the cancellation of the specified migrations.

* - 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 command stops the migration, and frees any resources it was using.
- The migration has not been committed. The command returns source and target devices or extents 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 the `dm migration move` command.

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`
- `dm migration pause`
- `dm migration remove`
- `dm migration resume`
- `dm migration start`

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
```

```
[-m|--migrations] context-path,context-path...
```

```
[-f|--force]
```

```
[-e|--rename-target]
```

Arguments

Required arguments	
<code>[-m --migrations]</code> <i>context-path,context-path...</i>	* List of one or more migrations to clean. Entries must be separated by commas.
Optional arguments	
<code>[-f --force]</code>	Forces the cancellation of the specified migrations.
<code>[-e --rename-target]</code>	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 devices 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.

Examples

```
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`
- `dm migration commit`
- `dm migration pause`
- `dm migration remove`
- `dm migration resume`
- `dm migration start`

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
```

```
[-m|--migrations] context-path,context-path...
```

```
[-f|--force]
```

Arguments

Required arguments	
--------------------	--

<code>[-m --migrations] context-path,context-path...</code>	* List of one or more migrations to commit. Entries must be separated by commas.
<code>[-f --force]</code>	Forces the commitment of the specified migrations.

* - 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.

Examples

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`
- `dm migration pause`
- `dm migration remove`
- `dm migration resume`
- `dm migration start`

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	
<code>[-m --migrations] context-path,context-path...</code>	* List of one or more migrations to pause. Entries must be separated by commas.

* - argument is positional.

Description

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`
- `dm migration commit`
- `dm migration remove`
- `dm migration resume`
- `dm migration start`

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`

`[-m|--migrations] context-path,context-path...`

`[-f|--force]`

Arguments

Required arguments	
--------------------	--

<code>[-m --migrations] context-path,context-path...</code>	* List of one or more migrations to remove. Entries must be separated by commas.
<code>[-f --force]</code>	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`
- `dm migration commit`
- `dm migration pause`
- `dm migration resume`
- `dm migration start`

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	
<code>[-m --migrations] context-path,context-path...</code>	List of one or more migrations to resume. Entries must be separated by commas.

Description

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
- dm migration commit
- dm migration pause
- dm migration remove
- dm migration start

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

```
[ -n | --name ] migration-name...
[ -f | --from ] { source-extent | source-device }
[ -t | --to ] { target-extent | target-device }
[ -s | --transfer-size ] value
--paused
--force
```

Arguments

Required arguments	
[-n --name] <i>migration-name...</i>	* Name of the new migration. Used to track the migration's progress, and to manage (cancel, commit, pause, resume) the migration.
[-f --from] { <i>source-extent</i> <i>source-device</i> }	* 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] { <i>target-extent</i> <i>target-device</i> }	* 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	
<code>[-s --transfer-size] <i>value</i></code>	<p>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.</p> <p>Range: 40 KB - 128 M. Default: 128 K.</p> <p>If the host I/O activity is very high, setting a large transfer size may impact host I/O. See About transfer-size in the batch-migrate start command.</p>
<code>--paused</code>	Starts the migration in a paused state.
<code>--force</code>	Do not ask for confirmation. Allows this command to be run using a non-interactive script.

* - 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.

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,
- De-fragment 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>
```

 **Note:** Extent migrations are blocked if the associated virtual volume is undergoing expansion. See the `virtual-volume expand` command.

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

This command can fail on a cross-cluster migration if there is not a sufficient number of meta volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

Start and manage device migrations from the device migration context:

```
VPlexcli:/> cd /data-migrations/device-migrations/
VPlexcli:/data-migrations/device-migrations>
```

When running the `dm migration start` command across clusters, you might receive the following error message:

```

VPlexcli:/> dm migration start -f SveTest_tgt_r0_case2_1_0002 -t
SveTest_src_r0_case2_2_0002 -n cc2
The source device 'SveTest_tgt_r0_case2_1_0002' has a volume
'SveTest_tgt_r0_case2_1_0002_vol' in a view. Migrating to device
'SveTest_src_r0_case2_2_0002' will create a synchronous distributed device.
In this GEO system, this can increase the per I/O latency on
'SveTest_tgt_r0_case2_1_0002_vol'. If applications using
'SveTest_tgt_r0_case2_1_0002_vol' are sensitive to this latency, they may
experience data unavailability. Do you wish to proceed ? (Yes/No) y
dm migration start: Evaluation of <<dm migration start -f
SveTest_tgt_r0_case2_1_0002 -t SveTest_src_r0_case2_2_0002 -n cc2>> failed.
cause: Failed to create a new data-migration.
cause: Unable to attach mirror
'SveTest_src_r0_case2_2_0002' to distributed Device 'MIGRATE_cc2'.
cause: Firmware command error.
cause: Active metadata device does not have a free
slot.

```

See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for instructions on increasing the number of slots.

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.

See the Dell EMC VPLEX Administration Guide for detailed information on data migration.

See also

- `batch-migrate create-plan`
- `batch-migrate start`
- `dm migration cancel`
- `dm migration commit`
- `dm migration pause`
- `dm migration remove`
- `dm migration resume`

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	
<code>[-v --storage-view] <i>context-path,context-path...</i></code>	List of one or more views to drill down. Entries must be separated by commas. Glob style pattern matching is supported.
<code>[-o --virtual-volume] <i>context-path,context-path...</i></code>	List of one or more virtual volumes to drill down. Entries must be separated by commas. Glob style pattern matching is supported.
<code>[-r --device] <i>context-path,context-path...</i></code>	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`.

Examples

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-FC00
```

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

ds dd convert-to-local

To convert a distributed device to a local device, this command detaches the leg that is not on the specified cluster.

Context

All contexts

Syntax

```
ds dd convert-to-local
[-h | --help]
[-v | --verbose]
[[-c | --cluster = ]cluster-context]
[-f | --force]
[[-d | --distributed-device=] distributed device]
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command. --verbose Provides more output during command execution. This may not have any effect for some commands.
-c --cluster= <i>cluster context</i>	Specifies the context path of the cluster where the distributed device will be local. If the device is exported to any cluster it must be the chosen cluster.

<code>-f --force</code>	Forces the command to proceed, bypassing all user warnings. positional arguments Specifies the distributed device to make local.
<code>-d --distributed-device= <i>context path</i></code>	Specifies the context path of the distributed device to make local.

The remaining leg becomes the supporting device of the virtual volume. The target device should NOT be migration temporary device and should not be exported to any other cluster than the specified cluster. For distributed devices that are part of a consistency-group please refer to the `consistency-group convert-local` command.

ds dd create

Creates a new distributed-device.

Contexts

All contexts.

Syntax

```
ds dd create
[-n|name] name
[-d|--devices] context-path [,contextpath,...]
[-l|--logging-volumes] context-path [,context-path,...]
[-r| rule-set] rule-set
[-s|--source-leg] context-path
[-f|--force]
```

Arguments

Required arguments	
<code>[-n --name] <i>name</i></code>	* The name of the new distributed device. Must be unique across the VPLEX.
<code>[-d --devices] <i>context-path</i> [, <i>context-path</i>,...]</code>	* List of one or more local devices that will be legs in the new distributed device.
<code>[-l --logging-volume] <i>context-path</i> [, <i>context-path</i>,...]</code>	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.
Optional arguments	
<code>[-r --rule-set] <i>rule-set</i></code>	The rule-set to apply to the new distributed device. If the <code>--rule-set</code> 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.

<code>[-s --source-leg]</code> <i>context-path</i>	Specifies one of the local devices to use 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.
<code>[-f --force]</code>	Forces a rule-set with a potential conflict to be applied to the new distributed device.

* - argument is positional.

Description

The new distributed device consists of two legs; local devices on each cluster.

⚠ WARNING Without `--source-leg`, 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. Do not use this command without `--source-leg` unless you plan to initialize the new device using host tools

⚠ 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 is only as large as the smaller local device
- The leftover capacity on the larger device is not 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 can contribute a maximum of one component device to the new distributed device.

This command can fail if there is not a sufficient number of meta volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

⚠ 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
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 it was halted. Only those portions of the device that have not been rebuilt are affected. The rebuild does not start over.

Examples

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/
cluster_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`
- `ds dd destroy`
- `local-device create`

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	
<code>[-c --cluster] <i>context-path</i></code>	* Specifies the winning cluster.
<code>[-d --distributed-device] <i>context-path</i></code>	Specifies the distributed device for which to declare a winning cluster.
<code>[-f --force]</code>	Forces the <code>declare-winner</code> 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.

Examples

```
VPlexcli:/distributed-storage/distributed-devices> ds dd declare-winner --
distributed-device DDtest_4 --cluster cluster-2 --force
```

See also

- `ds dd create`

ds dd destroy

Destroys the specified distributed-device(s).

Contexts

All contexts.

Syntax

```
ds dd destroy
[-d|--distributed-device] context-path, context-path,...
[-f|--force]
```

Arguments

Required arguments	
<code>[-d --distributed-device] <i>context-path, context-path,...</i></code>	* List of one or more distributed devices to destroy.
<code>[-f --force]</code>	Forces the distributed devices to be destroyed.

* - argument is positional.

Description

In order to be destroyed, the target distributed device must not host virtual volumes.

Examples

```
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

ds dd remove-all-rules

Removes all rules from all distributed devices.

Contexts

All contexts.

Syntax


```
ds dd remove-all-rules
[
-f|--force]
```

Arguments

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.

Examples

```
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
- ds rule island-containing
- ds rule-set copy
- ds rule-set create

- ds rule-set destroy
- ds rule-set what-if

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...*

[-c|--distributed-device-component] *context-path*

[-l|--logging-volumes] *context-path,context-path...*

[-n|--cancel]

Arguments

Required arguments	
[-d --distributed-devices] <i>context-path,context-path...</i>	One or more distributed devices for which segments of the specified logging volume are allocated/unallocated. All components of the distributed-device are included.
[-c --distributed-device-component] <i>context-path</i>	The distributed device component for which segments of the specified logging volume are allocated/unallocated.
[-l --logging-volume] <i>context-path,context-path...</i>	One or more logging-volumes where the logging volume segments are allocated/unallocated. The target volume must be created as <i>logging-volume</i> . 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.  WARNING Use the <code>--cancel</code> argument very carefully. the command issues a warning message if you attempt to cancel logging volumes on a distributed device. Removing the logging-volume for a device deletes the existing logging entries for that device. A FULL rebuild of the device occurs after a link failure and recovery. Removing the logging volume for all distributed devices removes 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.

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 occurs when the inter-cluster link is restored after an outage.
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.

This command can fail if there is not a sufficient number of meta volume slots. See the troubleshooting procedures for VPLEX in the SolVe Desktop for a resolution to this problem.

Examples

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
```

Attempt to cancel a logging volume for a distributed device that is not fully logged:

⚠ WARNING Issuing the cancel command on a distributed device that is not fully logged results in a warning message.

```
Vplexcli:/distributed-storage/distributed-devices/dr1_C12_0249> ds dd set-log --distributed-devices dr1_C12_0249 --cancel
WARNING: This command will remove the logging segments from distributed device 'dr1_C12_0249'.
If a distributed device is not fully logged, it is vulnerable to full rebuilds following
inter-cluster WAN link failure or cluster failure.
It is recommended that the removed logging-segments be restored as soon as possible.
```

See also

- logging-volume create

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	
<code>[-r --rule] <i>rule</i></code>	Specifies the rule to destroy.

Description

A rule-set contains rules. Use the `ll` command in the rule-set context to display the rules in the rule-set.

Examples

Use the `ds rule destroy` command to destroy a rule in the rule set.

```
VPlexcli:/distributed-storage/rule-sets/ruleset_recreate5/rules> ll
RuleName RuleType Clusters ClusterCount Delay Relevant
-----
rule_1 island-containing cluster-1 2 10s true
VPlexcli:/distributed-storage/rule-sets/ruleset_recreate5/rules> rule destroy
rule_1
```

See also

- `ds rule island-containing`
- `ds rule-set copy`
- `ds rule-set create`
- `ds rule-set destroy`
- `ds rule-set what-if`

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	
<code>[-c --clusters]</code> <i>context-path, context-path...</i>	* Clusters to which this rule applies.
<code>[-d --delay]</code> <i>delay</i>	* 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: min - delay timer in seconds s - delay timer in seconds (default unit) h - delay timer in hours
<code>[-r --rule-set]</code>	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:

1. VPLEX waits for 10 seconds after a link failure and then:
2. Resumes I/O to the island containing cluster-1,
3. 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`
- `ds rule destroy`
- `ds rule-set copy`
- `ds rule-set create`
- `ds rule-set destroy`
- `ds rule-set what-if`

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] new-rule-set
```

Arguments

Required arguments	
<code>[-s --source] <i>rule-set</i></code>	* Source rule-set.
<code>[-d --destination] <i>new-rule-set</i></code>	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
- ds rule destroy
- ds rule island-containing
- ds rule-set create
- ds rule-set destroy
- ds rule-set what-if

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	
<code>[-n --name] <i>rule-set</i></code>	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`
- `ds rule destroy`
- `ds rule island-containing`
- `ds rule-set copy`
- `ds rule-set create`
- `ds rule-set destroy`
- `ds rule-set what-if`
- `set`

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	
<code>[-r --rule-set] <i>rule-set</i></code>	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.

Examples

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`
- `ds rule destroy`
- `ds rule island-containing`
- `ds rule-set copy`
- `ds rule-set create`

- `ds rule-set what-if`
- `set`

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	
<code>[-i --islands] "<i>cluster-1,cluster-2</i>"</code>	List of islands, in quotes. The clusters are separated by commas, the islands by a space.
<code>[-r --rule-set] <i>context-path</i></code>	Context path of the rule-set used in the what-if scenario.

Description

This command supports only two clusters and one island.

Examples

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`
- `ds rule destroy`
- `ds rule island-containing`
- `ds rule-set copy`
- `ds rule-set create`
- `ds rule-set destroy`

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.

Table 15 ds summary field descriptions

Field	Description
Distributed Volumes (not in Consistency Groups) Unhealthy Summary:	
Device Name	Name of the device.
Health State	<ul style="list-style-type: none"> <code>major failure</code> - One or more children of the distributed device is out-of-date and will never rebuild, possibly because they are dead or unavailable. <code>minor failure</code> - 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. <code>non-recoverable error</code> - VPLEX cannot determine the distributed device's Health state. <code>ok</code> - The distributed device is functioning normally. <code>unknown</code> - VPLEX cannot determine the device's health state, or the state is invalid.
Operational Status	<ul style="list-style-type: none"> <code>degraded</code> - The distributed device may have one or more out-of-date children that will eventually rebuild.

Table 15 ds summary field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>error</code> - One or more components of the distributed device is hardware-dead. • <code>ok</code> - The distributed device is functioning normally. • <code>starting</code> - The distributed device is not yet ready. • <code>stressed</code> - One or more children of the distributed device is out-of-date and will never rebuild. • <code>unknown</code> - VPLEX cannot determine the distributed device's Operational state, or the state is invalid.
Service Status	<p><code>cluster unreachable</code> - VPLEX cannot reach the cluster; the status is unknown.</p> <p><code>need resume</code> - 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><code>need winner</code> - 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><code>potential conflict</code> - The clusters have detached each other resulting in a potential for detach conflict.</p> <p><code>running</code> - The distributed device is accepting I/O.</p> <p><code>suspended</code> - The distributed device is not accepting new I/O; pending I/O requests are frozen.</p> <p><code>winner-running</code> - 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.
Distributed volumes (in consistency groups) unhealthy summary:	

Table 15 ds summary field descriptions (continued)

Field	Description
CG Name	Name of the consistency group of which the unhealthy device is a member.
Cache Mode	Cache mode of the consistency group. Synchronous - 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.
Number of unhealthy volumes	Number of unhealthy volumes in the consistency group.
Cluster	Visibility of the consistency group. <ul style="list-style-type: none"> cluster-1 - consistency group is visible only at cluster-1. cluster-2- consistency group is visible only at cluster-2. cluster-1, cluster-2 - consistency group is visible at both clusters.
Operational Status	Current status for this consistency group with respect to each cluster on which it is visible. <ul style="list-style-type: none"> ok - I/O can be serviced on the volumes in the consistency group. suspended - I/O is suspended for the volumes in the consistency group. The reasons are described in the operational status: details. degraded - I/O is continuing, but there are other problems described in operational status: details. unknown - The status is unknown, likely because of lost management connectivity.
Status Details	If operational status is ok this field is empty: "[]". Otherwise, it displays additional information, which may be any of the following: <ul style="list-style-type: none"> requires-resolve-conflicting-detach - After the inter-cluster link is restored, two clusters have discovered that they have detached one another and resumed I/O

Table 15 ds summary field descriptions (continued)

Field	Description
	<p>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> <ul style="list-style-type: none"> • rebuilding-across-clusters - 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 <code>rebuild status</code> command to display which volume is out of date at which cluster. • rebuilding-within-cluster - One or more local rebuilds is in progress at this cluster. • data-safe-failure - 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. • requires-resume-after-data-loss-failure - 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 <code>consistency-group resume-after-data-loss-failure</code> command to select a winning cluster and allow I/O to resume. • cluster-departure - Not all the visible clusters are in communication. • requires-resume-after-rollback - A cluster has detached its peer cluster and rolled back the view of data, but is awaiting the <code>consistency-group resume-after-rollback</code> command before resuming I/O. Displayed: <ul style="list-style-type: none"> ▪ At the winning side when a detach rule fires, or shortly after the <code>consistency-group choose-winner</code> command picks a winning cluster.

Table 15 ds summary field descriptions (continued)

Field	Description
	<p><code>requires-resume-at-loser</code> - 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 <code>consistency-group resume-at-loser</code> command to make the view of data consistent with the winner, and to resume I/O at the loser.</p> <p><code>restore-link-or-choose-winner</code> - 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"> ■ There is no <code>detach-rule</code> ■ If the <code>detach-rule</code> is <code>'no-automatic-winner'</code>, or ■ If the <code>detach-rule</code> cannot fire because its conditions are not met. For example, if more than one cluster is active at the time of an inter-cluster link outage, the <code>'active-cluster-wins'</code> 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 <code>consistency-group choose-winner</code> command. <p><code>unhealthy-devices</code> - I/O has stopped in this consistency group because one or more volumes is unhealthy and cannot perform I/O.</p> <p><code>will-rollback-on-link-down</code> - 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>

Examples

Display summary information when no devices are unhealthy:

```
Vplexcli:/distributed-storage> ds summary
Slot usage summary:
```

```

    Total 912 slots used by distributed device logging segments.
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

```

Display summary information when one or more devices are unhealthy:

```

VPlexcli:/> ds summary
Slot usage summary:
    Total 912 slots used by distributed device logging segments.
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          Cluster    Operational
Status Details
-----
Unhealthy
Vols
-----
AA_ACW_Cluster12  synchronous  9                  cluster-1  unknown   []
                    [cluster-departure,
restore-link-or-choose-winner]
AP_ACW_Cluster1   synchronous  10                 cluster-1  unknown   []
                    [cluster-departure,
restore-link-or-choose-winner]
AP_ACW_Cluster2   synchronous  5                  cluster-1  unknown   []
                    [cluster-departure,
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.

```

Use the `--verbose` argument to display detailed information about unhealthy volumes in each consistency group:

```

VPlexcli:/> ds summary --verbose
Slot usage summary:
  Total 912 slots used by distributed device logging segments.
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 Cluster Operational
Status Details
-----
Unhealthy
-----
Vols
-----
AA_ACW_Cluster12 synchronous 9 cluster-1 unknown []
cluster-2 suspended
[cluster-departure,
restore-link-or-choose-winner]
AP_ACW_Cluster1 synchronous 10 cluster-1 unknown []
cluster-2 suspended
[cluster-departure,
restore-link-or-choose-winner]
AP_ACW_Cluster2 synchronous 5 cluster-1 unknown []
cluster-2 suspended
[cluster-departure,
restore-link-or-choose-winner]
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.
  luster cluster-1 : 25 distributed devices.
.

```

See also

- export port summary

- export storage-view summary
- extent summary
- local-device summary
- storage-volume summary
- virtual-volume provision

esrs import-certificate

Fetches the security certificate from a remote system and imports it to the local keystore. It resets the CLI process.

Contexts

All contexts.

Syntax

```
esrs import-certificate
-h | --help
--verbose
-f | --force
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.
-f --force	Bypasses input if none is provided.

Description

This command runs a interview script to import the SRSv3 security certificate.

 **Note:** In Metro systems, run this command on both the management servers.

Examples

Import ESRS certificate.

```
VPlexcli:/> esrs import-certificate
This command restarts the VPLEX CLI if the security settings are modified.
You will need to log in from all connected CLI and GUI
clients again.

Type CONTINUE to proceed, or press Enter to abort: CONTINUE

Please enter the IP address of the ESRS server: 10.228.67.5

-----Certificate Details-----

Owner: CN=esrshost, OU=ESRS, O=EMC, L=SO, ST=MA, C=US
Issuer: CN=esrshost, OU=ESRS, O=EMC, L=SO, ST=MA, C=US
Serial number: af0efd6abb77987f
Valid from: Thu Oct 25 05:44:37 UTC 2018 until: Mon Oct 25 05:44:37 UTC 2038
```

```

Certificate fingerprints:
  MD5:  C4:E6:AB:51:D0:0C:6C:0C:99:98:25:CC:75:3C:09:72
  SHA1: 8B:5D:EE:71:0B:38:DB:57:A3:B6:F2:DE:8E:71:0E:97:BB:10:EF:27
  SHA256:
B5:B6:8A:0B:FE:D8:0C:D9:CF:E3:9A:E3:3D:CD:81:6E:74:7C:CA:72:ED:0B:06:04:8E:F8:
24:53:B2:E6:6E:A5
Signature algorithm name: SHA256withRSA
Subject Public Key Algorithm: 2048-bit RSA key
Version: 3

Extensions:

#1: ObjectID: 2.5.29.19 Criticality=false
BasicConstraints:[
  CA:false
  PathLen: undefined
]

#2: ObjectID: 2.5.29.15 Criticality=false
KeyUsage [
  DigitalSignature
  Non_repudiation
  Key_Encipherment
]

Trust this certificate? (Y/N): Y
    
```

See also

- esrs register
- esrs un-register
- esrs status

esrs register

Registers the SRSv3 on the cluster.

Contexts

All contexts.

Syntax

```

esrs register
[-u|--username] user name
[-t|--password] {type
[-i|--ip-address] {ip_address
-h | --help
--verbose
    
```

Arguments

Required arguments	
[-u --username]	* Username of the user authorized to configure SRSv3 gateway on the VPLEX cluster.
[-p --password]	* The password to authenticate the user.

[-i --ip-address]	* The IP address of the SRSv3 gateway.
Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Adds information about a SRSv3 gateway to VPLEX. Used by VPLEX for enabling License Usage Data Transfer.

 **Note:** In Metro systems, run this command on both the management servers.

Examples

Register ESRS.

```
Vplexcli:/> esrs register -u john.norton@honeydew.com -p Password1 -i
10.228.67.5
ESRS successfully registered
```

See also

- esrs un-register
- esrs status
- esrs import-certificate

esrs status

Displays the SRSv3 status on the cluster.

Contexts

All contexts.

Syntax

```
esrs status
-h | --help
--verbose
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.

Description

Displays the connectivity status between VPLEX and SRSv3.

Note: In Metro systems, run this command on both the management servers.

Examples

ESRS status.

```
VPlexcli:/> esrs status

ESRS Gateway IP      Status      Cluster Name
-----
10.228.67.5         Online     cluster-1
```

See also

- esrs register
- esrs un-register
- esrs import-certificate

esrs un-register

Unregister the SRSv3 on the cluster.

Contexts

All contexts.

Syntax

```
esrs un-register
-h | --help
--verbose
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.

Description

Removes information about SRSv3 gateway from VPLEX.

Note: In Metro systems, run this command on both the management servers.

Examples

Unregister ESRS.

```
VPlexcli:/> esrs un-register

ESRS instance for this VPLEX(FNM00121500305) cluster un-registered
successfully.
```

See also

- `esrs register`
- `esrs status`
- `esrs import-certificate`

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] level`

`[-o|--component] component`

`[-m|--message] "message"`

Arguments

Required arguments	
<code>[-l --level] <i>level</i></code>	<p>Level of the event. Must be one of the following:</p> <p><code>emergency</code> - System is unusable.</p> <p><code>alert</code> - Immediate action is required.</p> <p><code>critical</code> - Critical condition detected.</p> <p><code>error</code> - Significant error condition detected.</p> <p><code>warning</code> - Warning condition is detected.</p> <p><code>notice</code> - Normal, but significant condition.</p> <p><code>info</code> - Information messages.</p> <p><code>debug</code> - Detailed event information used by Dell EMC for debugging.</p>
Optional arguments	
<code>[-d --directors] <i>context-path, context-path...</i></code>	One or more directors from which to send the test event.
<code>[-c --clusters] <i>context-path,context-path...</i></code>	One or more clusters from which to send the test event. Events are sent from every director in the specified cluster(s).
<code>[-o --component]</code>	Text to include in the component portion of the test message.
<code>[-m --message] "<i>message</i>"</code>	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 more directors 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.

All contexts.

Syntax

`exec` *command*

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

Optional arguments	
<code>[-e --exit-code] <i>exit-code</i></code>	Returns the specified value when the shell exits. If no exit code is specified, then 0 is returned.
<code>[-s --shutdown]</code>	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] seconds
[-w|--wait] seconds
[-c|--cluster] context-path
```

Arguments

Optional arguments	
<code>[-t --timeout] <i>seconds</i></code>	The maximum number of seconds to wait for the front-end fabric discovery operation to complete. Default: 300. Range: 1-3600.
<code>[-w --wait] <i>seconds</i></code>	The maximum number of seconds to wait for a response from the fabric discovery. Default: 10. Range: 1- 3600.
<code>[-c --cluster] <i>context-path</i></code>	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 might 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`

export initiator-port register

Registers an initiator-port, associating it with a SCSI address.

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] {type}
[-i|--initiator-port] initiator-port
[-p|--port] port

```


Arguments

Required arguments	
<code>[-i --initiator-port] <i>initiator-port</i></code>	* Name to assign to the registered port. Name must be unique in the system. Command fails if the specified name is already in use.
<code>[-p --port] <i>port</i></code>	* Port identifier. For Fibre Channel initiators, a WWN pair as follows: <i>portWWN nodeWWN</i> . <i>nodeWWN</i> is optional. Each WWN is either '0x' followed by one or more hex digits, or an abbreviation using the format: <i>string:number[,number]</i> . Following are four examples: <pre>0xd1342a 0xd1342b hyy1:194e,4 hyy1:194e 0xd1342a hyy1:194e,4</pre>
Optional arguments	
<code>[-c --cluster] <i>context-path</i></code>	Cluster on which the initiator port is registered.
<code>[-t --type] {<i>type</i>}</code>	Type of initiator port. If no type is specified, the default value is used. <ul style="list-style-type: none"> • <code>hpux</code> - Hewlett Packard UX • <code>sun-vcs</code> - Sun Solaris • <code>aix</code> - IBM AIX • <code>recoverpoint</code> - Dell EMC RecoverPoint • <code>ibm-d910</code> - IBM Series D910 • <code>default</code> - If no type is specified.

* - argument is positional.

Description

Use the `ll` command in `/engines/engine/directors/director/hardware /ports/ port` context to display portWWNs and nodeWWNs.

Registers an initiator-port and associates it with a SCSI address. For Fibre Channel, the SCSI address is represented by a WWN pair.

See also

- `export initiator-port discovery`
- `export initiator-port unregister`
- `export target-port renamewwns`
- `set`

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	
<code>[-f --file] file</code>	* The host declaration file path name.
Optional arguments	
<code>[-c --cluster] cluster-context</code>	* The cluster at which to create the view.
<code>[-p --ports] port,port...</code>	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.

See also

- `export initiator-port discovery`
- `export initiator-port unregister`

export initiator-port show-logins

Displays the initiator port logins to the VPLEX front-end target ports.

Context

All contexts.

Syntax

```
export initiator-port show-logins
[-i|--initiator-ports] initiator-ports[, initiator-ports...]
[-h|--help]
```

Arguments

Optional arguments	
<code>[-i --initiator-ports] <i>initiator-ports</i> [, <i>initiator-ports...</i>]</code>	* Specifies the initiator-ports for which the login information is required.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

Displays a list of target port logins for the specified initiator ports.

Example 2 Example

Shows target port logins for all the initiator ports in VPLEX:

```
Vplexcli:/> export initiator-port show-logins *
Cluster Names      Initiator Port Names      Target Port Names
-----
cluster-1          initiator_11               P000000003CA0014C-
A0-FC01
cluster-1          initiator_12               P000000003CA0014C-
A0-FC00
B0-FC00
cluster-2          initiator_21               P000000003CA00150-
A0-FC01
B0-FC01
cluster-2          initiator_22               P000000003CA00150-
A0-FC00
B0-FC00
P000000003CB00150-
```

Shows target port logins for initiator ports 11 and 22:

```
Vplexcli:/> export initiator-port show-logins -i
initiator_11,initiator_22
Cluster Names      Initiator Port Names      Target Port Names
-----
cluster-1          initiator_11               P000000003CA0014C-
A0-FC01
cluster-2          initiator_22               P000000003CA00150-
A0-FC00
B0-FC00
P000000003CB00150-
```

Shows no target port logins for initiator port 22:

```
Vplexcli:/> export initiator-port show-logins initiator_22
Cluster Names      Initiator Port Names      Target Port Names
```

Example 2 Example (continued)

```
-----
-----
cluster-2      initiator_22      None
```

See also

- `export initiator-port discovery`
- `export initiator-port register`
- `export initiator-port register-host`
- `export initiator-port unregister`

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
i|--initiator-port] initiator-port [,initiator-port...]
[-f|--force]
[-
```

Arguments

Required arguments	
<code>[-i --initiator-port] <i>initiator-port</i> [, <i>initiator-port</i>...]</code>	* One or more initiator ports to remove. Entries must be separated by commas.
Optional arguments	
<code>[-f --force]</code>	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`

export port summary

Displays a summary of exported ports for one or more clusters.

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,...]
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --clusters] cluster[, cluster,...]</code>	Display unhealthy ports for only the specified cluster(s).
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Displays the names of the unhealthy volumes exported on each port.

Description

Prints a summary of the views and volumes exported on each port, and a detailed summary of the unhealthy ports.

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.
VPLEXcli:/> export port summary -c cluster-1/
Port health summary(cluster-1):
port name                export status  view summary
-----
P0000000037204D11-ETH06 pending          view name                unhealthy
volumes operational status
-----
0                        degraded          chimera_view2_C1
0                        degraded          chimera_view1_C1
0                        degraded          C1_H2
0                        degraded          chimera_setupTearDown_C1
port name                health state  enabled  views  virtual-volumes
-----
P0000000037204D11-ETH07 healthy        true     4     30
P0000000037204D21-ETH06 healthy        true     4     30
P0000000037204D21-ETH07 healthy        true     4     30
P0000000037304D11-ETH06 healthy        true     4     30
P0000000037304D11-ETH07 healthy        true     4     30
P0000000037304D21-ETH06 healthy        true     4     30
P0000000037304D21-ETH07 healthy        true     4     30
P0000000037204D11-ETH06 unhealthy       false    4     30
    Total 8 ports, 1 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
- export storage-view summary
- extent summary
- local-device summary
- storage-volume summary
- virtual-volume provision

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] initiator-ports [, initiator-ports...]
```

Arguments

Required arguments	
<code>[-i --initiator-ports] initiator-ports [, initiator-ports...]</code>	* List of one or more initiator ports to add to the view. Entries must be separated by commas.
Optional arguments	
<code>[-v --view] context-path</code>	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`
- `export storage-view removeinitiatorport`

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	
<code>[-p --ports] <i>context-path,context-path ...</i></code>	* List of one or more ports to be added to the view. Entries must be separated by commas.
Optional arguments	
<code>[-v --view] <i>context-path</i></code>	Storage view to which to add the specified ports.

* - 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`
- `export storage-view removeport`

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
[ -o | --virtual-volumes ] virtual-volume, virtual-volume...
[ -f | --force ]
```

Arguments

Required arguments	
<code>[-o --virtual-volumes] <i>virtual-volume, virtual-volume ...</i></code>	<p>* List of one or more virtual volumes or LUN-virtual-volume pairs. Entries must be separated by commas.</p> <p>LUN-virtual-volume pairs must be enclosed in parentheses (). Virtual volumes and LUN-virtual-volume pairs can be typed on the same command line.</p> <p>When only virtual volumes are specified, the next available LUN is automatically assigned by VPLEX.</p>
Optional arguments	
<code>[-v --view] <i>context-path</i></code>	View to add the specified virtual volumes to.
<code>[-f --force]</code>	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.

 **Note:** You cannot add a virtual volume to a storage view if the initialization status of the virtual volume is `failed` or `in-progress`.

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
```

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)]
```

Add a virtual volume to a view using the `--force` option from the root context:

```
VPlexcli:/> export storage-view addvirtualvolume --view /clusters/Sau11/
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`
- `export storage-view create`
- `export storage-view removevirtualvolume`

- `virtual-volume create`
- `virtual-volume re-initialize`

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`
- `export storage-view find`
- `export storage-view map`
- `export storage-view show-powerpath-interfaces`

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	
<code>[-n --name] <i>name</i></code>	* Name of the new view. Must be unique throughout VPLEX.
<code>[-p --ports] <i>context-path,context-path...</i></code>	* List of one or more ports to add to the view.
Optional arguments	
<code>[-c --cluster] <i>context-path</i></code>	The cluster to create the view on.

* - 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.

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 front-end 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 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 front-end 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
- export storage-view addinitiatorport
- export storage-view addvirtualvolume
- export storage-view destroy

export storage-view destroy

Destroys the specified storage view.

Contexts

All contexts.

Syntax

```
export storage-view destroy
```

```
[-v|--view] context-path
```

```
[-f|--force]
```

Arguments

Required arguments	
<code>[-v --view] <i>context-path</i> ...</code>	* Storage view to destroy.
Optional arguments	
<code>[-f --force]</code>	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
- export storage-view removeinitiatorport
- export storage-view removeport

- `export storage-view removevirtualvolume`

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

Optional arguments	
<code>[-c --cluster] <i>cluster</i></code>	Cluster to search for views.
<code>[-v --volume] <i>volume</i></code>	Find the views exporting the specified volume. Identify the volume by name, VPD83 identifier, or a name pattern with wildcards.
<code>[-l --lun] <i>LUN</i></code>	Find the views exporting the specified LUN number.
<code>[-i --initiator-port] <i>initiator</i></code>	Find the views including the specified initiator. May contain wildcards.
<code>[-f --free-lun] -</code>	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`
- `export storage-view find-unmapped-volumes`
- `export storage-view map`
- `export storage-view summary`

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	
<code>[-c --cluster] <i>cluster</i></code>	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 unexported.

See also

- `export storage-view addvirtualvolume`
- `export-storage-view removevirtualvolume`

export storage-view map

Displays all the virtual volumes that are exported to the storage view.

Contexts

All contexts.

In `/clusters/cluster/exports` context, command is `storage-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	
<code>[-v --views] <i>view,view...</i></code>	* List of one or more storage views to map. Entries must be separated by commas. May contain wildcards.
Optional arguments	
<code>[-f --file] <i>file</i></code>	Name of the file to send the output to. 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`
- `export storage-view find`
- `export storage-view summary`

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	
<code>[-i --initiator-ports] <i>context-path,context-path...</i></code>	* Comma- separated list of one or more initiator ports to remove.
Optional arguments	
<code>[-v --view] <i>context-path</i></code>	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
- export storage-view removeport

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..`

Syntax

```
export storage-view removeport
[-v|--view] context-path
[-p|--ports] context-path,context-path...
```

Arguments

Required arguments	
<code>[-p --ports] context-path,context-path ...</code>	* List of one or more ports to be removed from the view. Entries must be separated by commas.
Optional arguments	
<code>[-v --view] context-path</code>	View from which to remove the specified ports.

* - 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
- export storage-view destroy

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
[o--virtual-volumes] volume,volume...
[-f|--force]
```

Arguments

Required arguments	
<code>[<i>o</i>--virtual-volumes]</code> <i>volume,volume ...</i>	* List of one or more virtual volumes to be removed from the view. Entries must be separated by commas.
Optional arguments	
<code>[-f --force]</code>	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.
<code>[-v --view] <i>context-path</i></code>	View from which to remove the specified virtual volumes.

* - 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`

- `export storage-view destroy`

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

Optional arguments	
<code>[-c --cluster] context-path</code>	The cluster at which to show the PowerPath interface mapping.

See also

- `export storage-view checkconfig`
- `export storage-view find`
- `export storage-view map`
- `export storage-view summary`

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

Optional arguments	
<code>[-c --cluster] cluster, cluster...</code>	List of clusters. Entries must be separated by commas. Display information only for storage views on the specified clusters.

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
-----
-----
esx1_view error                port name                unhealthy volumes
export status
-----
ok                P000000003B2017D8-A0-FC00 1
ok                P000000003B2017D8-A0-FC01 1
ok                P000000003B3017D8-B0-FC00 1
ok                P000000003B3017D8-B0-FC01 1
view name  health-state  exported volumes  port s  registered initiators
-----
esx1_view unhealthy        10                4      2
          Total 1 views, 1 unhealthy.
```

See also

- `export port summary`
- `export storage-view checkconfig`
- `export storage-view map`
- `export storage-view show-powerpath-interfaces`
- `storage-volume summary`

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	
<code>[-w --wwns] <i>wwns</i></code>	<p>A WWN pair separated by “ ”: <i>portWWN nodeWWN</i></p> <p>Each WWN is either '0x' followed by one or more hexadecimal digits or an abbreviation, in the following format: <i>string.number[,number]</i></p> <p>For example, 0xd1342a 0xd1342b hyy1:194e,4 hyy1:194e 0xd1342a hyy1:194e,4</p>
Optional arguments	
<code>[-p --port] <i>context-path</i></code>	- 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
P00000000000000001-FK00
```

See also

- export initiator-port discovery

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] <i>storage-volume,storage-volume ...</i>	* Names of one or more claimed storage volumes to extent. Entries must be separated by commas.
Optional arguments	
[-s --size] <i>size</i>	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] <i>integer</i>	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] <i>integer</i>	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 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 virtual volume, create a single extent per storage volume, and then use devices to concatenate or stripe these extents into a larger device.

This command can fail if there is not a sufficient number of meta volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

Examples

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
Basic_cl_ramdisk_100GB_684	VPD83T3:60001440000000103017dfea88355431	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_685	VPD83T3:60001440000000103017dfea88355433	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_686	VPD83T3:60001440000000103017dfea88355435	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_687	VPD83T3:60001440000000103017dfea88355437	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_688	VPD83T3:60001440000000103017dfea88355439	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_689	VPD83T3:60001440000000103017dfea8835543b	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_68	VPD83T3:60001440000000103017dfea88354f61	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_690	VPD83T3:60001440000000103017dfea8835543d	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_691	VPD83T3:60001440000000103017dfea8835543f	100G	claimed	EMC	alive	normal
Basic_cl_ramdisk_100GB_692	VPD83T3:60001440000000103017dfea88355441	100G	claimed	EMC	alive	normal

```
VPlexcli:/>cd /clusters/cluster-1/storage-elements/storage-volumes
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> extent create
Symm1723_1DC,Symm1723_1E0
```

The following example shows creating two extents on top of a thin-capable storage volume (with the restriction that a thick extent will be created):

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> extent create
myVolume --num-extents 2
You are creating 2 extents on top of 1 thin-capable storage-volume
'myVolume'. The resulting extents will not
be thin-capable.

VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes>
```

The following example shows creating an extent that is smaller than the supporting storage volume (with the restriction that a thick extent will be created):

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> extent create
myVolume --size 1MB
The new extent will not completely encompass the following thin-capable
storage-volume: myVolume. The resulting extent will not be thin-capable.

VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes>
```

See also

- extent create
- extent destroy

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	
--------------------	--

<code>[-s --extents] context-path,context-path ...</code>	* List of one or more extents to destroy. Entries must be separated by commas.
Optional arguments	
<code>[-f --force]</code>	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`

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

Optional arguments	
<code>[-c --clusters] cluster,cluster...</code>	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.

Table 16 extent summary field descriptions

Field	Description
Health summary (displayed only for unhealthy extents)	
Name	Name of extent.
I/O Status	<p>alive - I/O is proceeding normally on the extent.</p> <p>dead - The underlying storage volume is marked as hardware-dead.</p> <p>unreachable - The underlying storage volume is unreachable.</p>
Operational Status	<p>degraded - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).</p> <p>ok - The extent is functioning normally.</p> <p>starting - The extent is not yet ready.</p> <p>unknown - VPLEX cannot determine the extent's Operational state, or the state is invalid.</p>
Health State	<p>degraded - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).</p> <p>ok - The extent is functioning normally.</p> <p>non-recoverable-error - 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.</p> <p>unknown - VPLEX cannot determine the extent's Operational state, or the state is invalid.</p>
Extent Summary	
Health	<p>extents - Total number of extents on the cluster, the number.</p> <p>out-of-date - Of the total number of extents on the cluster, the number that are out-of-date compared to their mirror.</p> <p>unhealthy - Of the total number of extents on the cluster, the number with operational status or health state that is not "ok".</p>
Use	<p>used - Of the total number of extents on the cluster, the number in use.</p> <p>claimed - Of the total number of extents on the cluster, the number that are claimed</p> <p>unclaimed - Of the total number of extents on the cluster, the number that are unclaimed.</p>

Table 16 extent summary field descriptions (continued)

Field	Description
	<p>unusable - 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.</p> <p>logging - Of the total number of extents on the cluster, the number that are in use for logging.</p>
Capacity	Total capacity on the cluster.

See also

- ds summary
- export port summary
- export storage-view summary
- local-device summary
- storage-volume summary
- virtual-volume provision

find

Finds all the contexts matching a pattern and returns a set contexts matching supplied pattern.

Contexts

All contexts.

Syntax

find

[-c | --contexts] = *pattern* [, *pattern* ...]

[-h | --help]

[--verbose]

Arguments

Required arguments	
[-c --contexts] = <i>pattern</i> [, <i>pattern</i> ...]	Pattern for matching contexts you want to find.
Optional arguments	
[-h --help]	Displays the usage for this command
[--verbose]	Provides additional output during command execution. This may not have any effect for some commands.

Description

Use this command to find all contexts matching a pattern. When invoked interactively, the command prints the contexts to the screen.

See [Searching the context tree](#) for more information about the find command and related examples.

front-end-performance-stats start

Starts the collection of the read and write statistics with the I/O size and the logical block addressing (LBA) information on the VPLEX virtual volumes through periodic polling.

Contexts

All contexts.

Syntax

```
front-end-performance-stats start
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This might not have any effect for some commands.

Description

Starts the collection of the read and write statistics with the I/O size and the logical block addressing (LBA) information on the VPLEX virtual volumes through periodic polling. This command starts generating the performance data, which helps resolve I/O performance issues with VPLEX. The statistics are available in the `fe_perf_stats_<timestamp>.log` file at `/var/log/VPlex/cli/`.

Note: Run this command on each cluster to collect the front-end performance statistics. After you run this command, the system continues to collect the front-end performance statistics until you run the `front-end-performance-stats stop` command.

See also

- `front-end-performance-stats stop`
- `front-end-performance-stats status`

front-end-performance-stats status

Displays the status of front-end performance statistics collection.

Contexts

All contexts.

Syntax

```
front-end-performance-stats status
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This might not have any effect for some commands.

Description

Provides the status of the front-end performance statistics collection. The details include the name of the log file that contains the statistics, the period for which the statistics collection has been running, the time when the directors were polled for information for the last time, and the number of errors that occurred per director in the last two hours.

Note: Run this command on each cluster to view the status of the front-end performance statistics collection.

See also

- `front-end-performance-stats start`
- `front-end-performance-stats stop`

front-end-performance-stats stop

Stops the front-end performance statistics collection.

Contexts

All contexts.

Syntax

```
front-end-performance-stats stop
```

Arguments

Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This might not have any effect for some commands.

Description

Stops the front-end performance statistics collection. The statistics are available in the `fe_perf_stats_<timestamp>.log` file at `/var/log/VPlex/cli/`.

Note: Run this command on each cluster to stop the front-end performance statistics collection.

See also

- `front-end-performance-stats start`
- `front-end-performance-stats status`

getsysinfo

Returns information about the current system.

Contexts

All contexts.

Syntax

```
getsysinfo
--output path-name
-linix
```

Arguments

Optional arguments	
--output <i>path-name</i>	Location and name of the output file. Default: /var/log/VPlex/cli/YY-sysinfo.txt
--linix	Use this if the management server is running on a Linux system. Disables the scsi tests since Linux systems lack a scsi command.

Description

Display information and send the output to a file.

The information is written in TCL format.

Table 17 getsysinfo field descriptions

Field	Description
Flag includeDebug	Ignore this line.
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.

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
```

See also

- cluster summary
- director firmware show-banks
- manifest version
- version

health-check

Displays a report indicating overall hardware/software health.

Contexts

All contexts.

Syntax

```
health-check
[-m|--highlevel]
[-f|--full]
--configuration
--back-end
--front-end
--limits
--cache
--consistency-group
--wan
--cluster_witness
```

Arguments

Optional arguments

<code>[-m --highlevel]</code>	Checks for major subcomponents with error conditions. Warnings are ignored. Used for instantaneous, high level view of the health of the VPLEX. Default behavior if no other argument is specified.
<code>[-f --full]</code>	Runs full scan.
<code>--configuration</code>	Perform configuration checks.
<code>--back-end</code>	Perform back end checks.
<code>--front-end</code>	Perform front end checks.
<code>--cache</code>	Perform cache check.
<code>--limits</code>	Lists the configuration limits for the configuration.
<code>--consistency-group</code>	Perform consistency group check.
<code>--wan</code>	Perform WAN health checks.
<code>--cluster_witness</code>	Perform cluster witness related checks.

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      Total    Unhealthy  Total    Total    Total    Total
Name         Storage Unhealthy Registered Ports    Exported ITLs
Views       Views   Initiators
-----
cluster-1    4        2          12       8        135     672
cluster-2    0        0          0        0        0       0
Storage:
-----
Cluster      Total    Unhealthy  Total    Unhealthy  Total  Unhealthy  No
Not visible Name     Storage Storage  Virtual  Virtual    Dist  Dist      Dual
from      Volumes Volumes  Volumes Volumes   Devs  Devs      Paths
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      Total    Unhealthy  Total    Total    Unhealthy
Name         Synchronous Synchronous Asynchronous Asynchronous
Groups      Groups      Groups      Groups
-----
cluster-1    9        0          0        0          0
cluster-2    5        0          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      Total    Unhealthy  Total    Unhealthy  Total    Mis-
aligned Total    Unhealthy  Total    Unhealthy  Total    Mis-
Name         RP Clusters RP Clusters Replicated Replicated RP-enabled RP-
enabled     Registered  RP Storage  Virtual  Virtual    Consistency
Consistency RP Initiators/ Views  Volumes  Volumes  Groups
Groups      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.
Array Aware:
-----
Cluster Name   Provider   Address           Connectivity   Registered   Total
-----
Arrays        Storage Pool
-----
Hopkinton     dsveal25  10.108.64.125    connected     2            13
Hopkinton     dsveal23  10.108.64.123    connected     2            29
Providence    dsveal24  10.108.64.124    connected     3            21
```

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
Checking COM Port Power Level..... OK
Checking Meta Data Backup..... Warning
Checking Meta Data Slot Usage..... Error
Back End (BE):
Checking Unreachable Storage Volumes..... OK
Checking Degraded Storage Volumes.....Error
Checking Unhealthy Virtual Volumes..... OK
Back end array status..... OK
Validating paths to back end arrays..... OK
Front End (FE):
Checking Storage Views..... OK
Checking Front End Path..... OK
Cache:
Checking for sub-pages writes(25% or above of total writes).... OK
Checking Stuck IO..... OK
Consistency Group Health:
Consistency Group Health..... OK
WAN Link:
WAN Configuration..... OK
WAN Port Settings..... OK
WAN Port Group Settings..... OK
WAN Subnet Settings..... OK
WAN Bridged and Routed Settings..... OK
WAN Ping Remote IPs..... OK
Director Health Status:
Checking SSD Hardware..... OK
Checking Director RPMs..... OK
Output to /var/log/VPlex/cli/health_check_full_scan.log
```

Run 2 WAN-only scans:

- The first in (default) high-level mode,
- The second in verbose mode:

```
VPlexcli:/> health-check --wan
WAN Link:
WAN Configuration.... OK
```

```

VPlexcli:/> health-check --wan --verbose
WAN Link:
WAN Configuration..... OK
Connectivity: full
    All port-groups have a status of either ok or warning. com connectivity is
    complete through minor configuration errors may still exist (see
individual
    port-group statuses).
port-group-1: OK
    All com links have the expected connectivity: this port-group is operating
    correctly.
port-group-0: OK
    All com links have the expected connectivity: this port-group is operating
    correctly.
WAN Port Settings..... OK
WAN Port Group Settings..... OK
WAN Subnet Settings..... OK
WAN Bridged and Routed Settings.... OK
WAN Ping Remote IPs..... OK
ping from director-1-1-A:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.328 ms
    Remote Cluster IPs:
        192.168.4.252 is reachable, Round-trip time:0.083 ms
        192.168.5.252 is reachable, Round-trip time:0.081 ms
ping from director-1-1-B:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.431 ms
    Remote Cluster IPs:
        192.168.4.252 is reachable, Round-trip time:0.086 ms
        192.168.5.252 is reachable, Round-trip time:0.083 ms
ping from director-1-2-A:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.297 ms
    Remote Cluster IPs:
        192.168.4.252 is reachable, Round-trip time:0.088 ms
        192.168.5.252 is reachable, Round-trip time:0.103 ms
ping from director-1-2-B:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.311 ms
    Remote Cluster IPs:
        192.168.4.252 is reachable, Round-trip time:0.089 ms
        192.168.5.252 is reachable, Round-trip time:0.075 ms
ping from director-2-1-A:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.371 ms
    Remote Cluster IPs:
        192.168.4.251 is reachable, Round-trip time:0.08 ms
        192.168.5.251 is reachable, Round-trip time:0.091 ms
ping from director-2-1-B:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.443 ms
    Remote Cluster IPs:
        192.168.4.251 is reachable, Round-trip time:0.066 ms
        192.168.5.251 is reachable, Round-trip time:0.083 ms
ping from director-2-2-A:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.347 ms
    Remote Cluster IPs:
        192.168.4.251 is reachable, Round-trip time:0.087 ms
        192.168.5.251 is reachable, Round-trip time:0.087 ms
ping from director-2-2-B:
    Remote Discovery IP:
        224.100.100.100 is reachable, Round-trip time:0.397 ms
    Remote Cluster IPs:
        192.168.4.251 is reachable, Round-trip time:0.088 ms

```

```
192.168.5.251 is reachable, Round-trip time:0.087 ms
Output to /var/log/VPLEX/cli/health_check_full_scan.log
```

Check the back end of a cluster

```
Vplexcli:/> health-check --back-end --verbose
Back End (BE):
Checking Unreachable Storage Volumes..... OK
Checking Degraded Storage Volumes..... Error
Degraded storage volumes found

Error

Cluster cluster-1:
  There are 8 storage volumes running in degraded mode.

  Array: EMC-CLARiiON-APM00114102495
    There are 8 storage volumes running in degraded mode.
    First 4 storage volumes in degraded mode are:
    VPD83T3:600601601dd028007a09dalb6427e111 is degraded ['degraded-
timeout', 'degraded-read-write-latencies']
    VPD83T3:600601601dd028007fc9ec0e6427e111 is degraded ['degraded-read-
write-latencies']
    VPD83T3:600601601dd0280080c9ec0e6427e111 is degraded ['degraded-
timeout', 'degraded-write-latency']
    VPD83T3:600601601dd0280083c9ec0e6427e111 is degraded ['degraded-write-
latency']

Checking Unhealthy Virtual Volumes..... OK
Back end array status..... OK
cluster-1 EMC-CLARiiON-APM00114102495 connectivity status is ok
cluster-1 EMC-SYMMETRIX-195700501 connectivity status is ok
cluster-1 SANBlaze-VLUNP5T0-110d connectivity status is ok
cluster-2 EMC-CLARiiON-APM00114102489 connectivity status is ok
cluster-2 EMC-SYMMETRIX-195700501 connectivity status is ok
cluster-2 SANBlaze-VLUNP13T0-110d connectivity status is ok

Validating paths to back end arrays..... OK

Output to /home/service/vafadm/cli/health_check_full_scan.log
```

Health-check limits on a Metro system.

```
Vplexcli:health-check --limits
Product Version: 6.1.1.00.00.04
Product Type: Metro
WAN Connectivity Type: FC
Hardware Type: VS2
Cluster Size: 2 engines
Cluster TLA:
  cluster-1: FNM00121500305
  cluster-2: FNM00121300045

Cluster Configuration Limits:
-----
Configuration                                                    Maximum
-----
Active intra-cluster rebuilds                                     25
Maximum WAN latency (RTT) in a VPLEX Metro                       5 ms
Local top-level devices                                          8000
Extents                                                           24000
Storage volumes                                                  8000
```

Virtual volume size	64TB
Storage views	1000
Clusters	2
Extents per storage volume	128
Volumes per Consistency Group	1000
Storage volume size	64TB
IT nexus per cluster in VPLEX	3200
Distributed devices - includes distributed and local devices with global visibility	8000
Total storage provisioned in a system	8PB
LUNs exported through each IT nexus on VPLEX back-end	4096
Extent block size	4 KB
Active inter-cluster rebuilds (on distributed devices)	25
IT nexus per back-end port	256
Synchronous Consistency Groups	1024
RAID 1 mirror legs	2
Virtual volumes	8000
Minimum bandwidth for VPLEX Metro IP WAN link	3 Gbps
IT nexus per front-end port	400
Paths per storage volume per VPLEX director	4
Minimum bandwidth for VPLEX Metro with RAPIDPath IP WAN link	1 Gbps

See also

- cluster status
- validate-system-configuration

help

Displays help on one or more commands.

Contexts

All contexts.

Syntax

help

[-i|--interactive]

[-G|--no-global]

[-n|--no-internal]

Arguments

Optional arguments	
[-i --interactive]	Invoke interactive help. Type <code>quit</code> to exit 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.

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

Optional arguments	
<code>[-c --clear]</code>	Clears the history list.
<code>[-n --number] <i>number</i></code>	Displays only the last <i>number</i> commands in the history list.

Example

Display the last 8 commands executed in this CLI session:


```
VPlexcli: /> 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
```

iscsi chap back-end add-credentials

Adds one or more configuration credentials for back-end CHAP.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap back-end add-credentials
```

```
[-u|--username] username
```

```
[-s|--secret] secret
```

```
[-t|--targets] target [, target...]
```

```
[-c|--cluster] cluster-context
```

```
[--secured]
```

```
[-h|--help]
```

```
[--verbose]
```

Arguments


Required arguments	
<code>[-u --username] <i>username</i></code>	* Specifies the username to use in the configured CHAP credentials.
<code>[-s --secret] <i>secret</i></code>	* Specifies the secret to use in the configured CHAP credentials. The secret must be between 12 and 255 characters long, and can be composed of all printable ASCII characters.

<code>[-t --targets] <i>target</i> [, <i>target...</i>]</code>	* Specifies the IQNs or IP addresses of the targets for which to configure credentials.
<code>[--secured]</code>	Prevents the secret from being stored as plain text in the log files.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context-path of the cluster at which the back-end CHAP credentials should be added.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Adds one or more configuration credentials for back-end CHAP.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Add back-end CHAP credentials to targets on cluster 1:

```
VPlexcli: /> iscsi back-end add-credentials user_name secret_code12345
10.10.100.1 -c /clusters/cluster-1 --secured
Successfully added credentials to 8 targets.
```

See also


- `iscsi chap back-end disable`
- `iscsi chap back-end enable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end disable

Disables back-end CHAP.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax


```
iscsi chap back-end disable
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```



Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which back-end CHAP should be disabled.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Disables back-end CHAP. Once CHAP is disabled on the back-end, all configured credentials will be unused, but not deleted.

 **WARNING** You will not be able to login to storage arrays that enforce CHAP if you disable back-end CHAP that was previously configured.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Disable back-end CHAP on cluster 1:

```
VPlexcli:/> iscsi chap back-end disable -c cluster-1/
Successfully disabled back-end chap.
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end enable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end enable

Enables back-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax


```
iscsi chap back-end enable
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which back-end CHAP should be enabled.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Enables back-end CHAP. Enabling CHAP on the back-end allows VPLEX to log in securely to storage arrays that also have CHAP configured. To use CHAP once it is enabled, all storage arrays must have added CHAP credentials or have set a default CHAP credential .

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Enable back-end CHAP on cluster 1:

```
VPlexcli:/> iscsi chap back-end enable -c cluster-1/
Successfully enabled back-end chap.
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end disable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end list-credentials

Lists all configured back-end CHAP credentials.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax


```
iscsi chap back-end list-credentials
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster for which the configured back-end CHAP credentials should be listed.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Lists all configured back-end CHAP credentials.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

List configured back-end CHAP credentials for cluster 1:

```
VPlexcli:/> iscsi chap back-end list-credentials -c cluster-1/
Credential Target Username      Secret
-----
10.10.100.1   user          chap_secret_123
10.10.100.2   vnx_username  vnx_secret_123456
all           default_username default_secret
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end disable`
- `iscsi chap back-end enable`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end remove-credentials

Removes one or more configuration credentials for back-end CHAP.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap back-end remove credentials
[-t|--targets] target [, target...]
[-c|--cluster] cluster-context
[-h|--help]
 [--verbose]
```


Arguments

Required arguments	
<code>[-t --targets] <i>target</i> [, <i>target...</i>]</code>	* Specifies the IQNs or IP addresses of the targets for which to remove credentials.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the back-end CHAP credential should be removed.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Removes one or more configuration credentials for back-end CHAP.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Remove a configured back-end CHAP credential on a specified IP address on cluster 1:

```
VPlexcli:/> iscsi chap back-end remove-credentials -c cluster-1/ 10.10.100.1
Successfully removed credentials to 8 targets.
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end disable`
- `iscsi chap back-end enable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end remove-default-credential

Removes a default credential for back-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap back-end remove-default-credential
[-c|--cluster] cluster-context
```

[-h|--help]


[--verbose]

Arguments

Optional arguments	
[-c --cluster] <i>cluster-context</i>	Specifies the context path of the cluster at which the default back-end CHAP credential should be removed.
[-h --help]	Displays command line help.
[--verbose]	Provides more output during command execution. This may not have any effect for some commands.

Description

Removes a default credential for back-end CHAP configuration.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Remove default credentials for back-end CHAP from ports on cluster 1:

```
VPlexcli:/> iscsi chap back-end remove-default-credential -c cluster-1/
Successfully removed default credentials on 8 ports.
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end disable`
- `iscsi chap back-end enable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end set-default-credential`

iscsi chap back-end set-default-credential

Sets a default credential for back-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap back-end set-default-credential
[-u|--username] username
[-s|--secret] secret
[-c|--cluster] cluster-context [--secured]
[-h|--help]
[--verbose]
```


Arguments

Required arguments	
<code>[-u --username] <i>username</i></code>	* Specifies the username to use in the configured CHAP credentials.
<code>[-s --secret] <i>secret</i></code>	* Specifies the secret to use in the configured CHAP credentials. The secret must be between 12 and 255 characters long, and can be composed of all printable ASCII characters.
<code> [--secured]</code>	Prevents the secret from being stored as plain text in the log files.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the default back-end CHAP credential should be set.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - positional argument

Description

Sets a default credential for back-end CHAP configuration. When back-end CHAP is enabled, VPLEX will log into its targets using the default credential. If a specific target has a separate credential configured, VPLEX will use that separate credential when attempting to log into the target.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Set a default credential for back-end CHAP configuration on cluster 1:

```
Vplexcli:/> iscsi chap back-end set-default-credential -c cluster-1/ -u
username -s secret123456 --secured
Successfully set default credential on 8 ports.
```

See also


- `iscsi chap back-end add-credentials`
- `iscsi chap back-end disable`
- `iscsi chap back-end enable`
- `iscsi chap back-end list-credentials`
- `iscsi chap back-end remove-credentials`
- `iscsi chap back-end remove-default-credential`

iscsi chap front-end add-credentials

Adds one or more configuration credentials for front-end CHAP.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap front-end add-credentials
[-u|--username] username
[-s|--secret] secret
[-i|--initiators] initiator [, initiator...]
[-c|--cluster] cluster-context
[--secured]
[-h|--help]
[--verbose]
```


Arguments

Required arguments	
<code>[-u --username] <i>username</i></code>	* Specifies the username to use in the configured CHAP credentials.
<code>[-s --secret] <i>secret</i></code>	* Specifies the secret to use in the configured CHAP credentials. The secret must be between 12 and 255 characters long, and can be composed of all printable ASCII characters.
<code>[-i --initiators] <i>initiator</i> [, <i>initiator</i>]</code>	* Specifies the IQNs of the initiators for which to configure credentials.
<code> [--secured]</code>	Prevents the secret from being stored as plain text in the log files.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the front-end CHAP credentials should be added.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Adds one or more configuration credentials for front-end CHAP.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Add configuration credentials for front-end CHAP to initiators on cluster 1:

```
Vplexcli:/> iscsi chap front-end add-credentials -c cluster-1/ -u user -s
secret123456 iqn.1991-05.com.microsoft:dahlia-w2k8 --secured
Successfully added credentials to 8 initiators.
```

See also

- `iscsi chap front-end disable`
- `iscsi chap front-end enable`
- `iscsi chap front-end list-credentials`
- `iscsi chap front-end remove-credentials`
- `iscsi chap front-end remove-default-credential`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end disable

Disables front-end CHAP.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```
iscsi chap front-end disable
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which front-end CHAP should be disabled.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Disables front-end CHAP. Once CHAP is disabled on the front-end all configured credentials will be unused, but not deleted.

Note: This command is valid only on systems that support iSCSI devices.

Example

Disable front-end CHAP on cluster 1:

```
Vplexcli: /> iscsi chap front-end disable -c cluster-1/
Successfully disabled front-end chap.
```


See also


- `iscsi chap front-end add-credentials`
- `iscsi chap front-end enable`
- `iscsi chap front-end list-credentials`
- `iscsi chap front-end remove-credentials`
- `iscsi chap front-end remove-default-credential`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end enable

Enables front-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax


```
iscsi chap front-end enable
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```


Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which front-end CHAP should be enabled.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Enables front-end CHAP. Enabling CHAP on the front-end will allow initiators to log in securely. To use CHAP once it is enabled, all initiators must have added CHAP credentials or have set a default CHAP credential .

 **WARNING** Initiators will no longer be able to log-in once CHAP is enabled if they do not have a credential configured.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Enable front-end CHAP on cluster 1:

```
Vplexcli: /> iscsi chap front-end enable -c cluster-1/
Successfully enabled front-end chap.
```

See also

- `iscsi chap front-end add-credentials`
- `iscsi chap front-end disable`
- `iscsi chap front-end list-credentials`
- `iscsi chap front-end remove-credentials`
- `iscsi chap front-end remove-default-credential`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end list-credentials

Lists all configured front-end CHAP credentials.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```
iscsi chap front-end list-credentials
[-c|--cluster] cluster-context [-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster for which the configured front-end CHAP credentials should be listed.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Lists all configured front-end CHAP credentials.

Note: This command is valid only on systems that support iSCSI devices.

Example

List all configured front-end CHAP credentials for cluster 1:

```
VPlexcli:/> iscsi chap front-end list-credentials -c cluster-1/
Credential Initiator                               Username Secret
-----
iqn.1991-05.com.microsoft:dahlia-w2k8 user        secret123456
```

See also

- `iscsi chap front-end add-credentials`


- `iscsi chap front-end disable`
- `iscsi chap front-end enable`
- `iscsi chap front-end remove-credentials`
- `iscsi chap front-end remove-default-credential`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end remove-credentials

Removes one or more configuration credentials for front-end CHAP.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap front-end remove-credentials
[-i|--initiators] initiator [, initiator]
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```


Arguments

Required arguments	
<code>[-i --initiators] <i>initiator</i> [, <i>initiator</i>]</code>	* Specifies the IQNs of the initiators for which to remove credentials.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the front-end CHAP credentials should be removed.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - positional argument

Description

Removes one or more configuration credentials for front-end CHAP.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Remove front-end CHAP credentials from initiators on cluster 1:

```
Vplexcli:/> iscsi chap front-end remove-credentials -c cluster-1/
iqn.1991-05.com.microsoft:dahlia-w2k8
Successfully removed credentials to 8 initiators.
```

See also


- `iscsi chap front-end add-credentials`
- `iscsi chap front-end disable`
- `iscsi chap front-end enable`
- `iscsi chap front-end list-credentials`
- `iscsi chap front-end remove-default-credential`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end remove-default-credential

Removes a default credential for front-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax


```
iscsi chap front-end remove-default-credential
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the default front-end CHAP credential should be removed.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Remove a default credential for front-end CHAP configuration.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Remove a default front-end CHAP credential configuration from cluster 1:

```
VPlexcli:/> iscsi chap front-end remove-default-credential -c cluster-1/
Successfully removed default credentials on 8 ports.
```

See also

- `iscsi chap front-end add-credentials`


- `iscsi chap front-end disable`
- `iscsi chap front-end enable`
- `iscsi chap front-end list-credentials`
- `iscsi chap front-end remove-credentials`
- `iscsi chap front-end set-default-credential`

iscsi chap front-end set-default-credential

Sets a default credential for front-end CHAP configuration.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi chap front-end set-default-credential
[-u|--username] username
[-s|--secret] secret
[-c|--cluster] cluster-context
[--secured]
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-u --username] <i>username</i></code>	* Specifies the username to use in the configured CHAP credentials.
<code>[-s --secret] <i>secret</i></code>	* Specifies the secret to use in the configured CHAP credentials. The secret must be between 12 and 255 characters long, and can be composed of all printable ASCII characters.
<code>[--secured]</code>	Prevents the secret from being stored as plain text in the log files.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which the default front-end CHAP credential should be set.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - positional argument

Description

Sets a default credential for front-end CHAP configuration. When front-end CHAP is enabled, any initiator without a separate credential will be able to login using this default credential. A specific initiator that has a separate credential configured will not be able to login using the default credential.

Note: This command is valid only on systems that support iSCSI devices.

Example

Set a default credential for front-end CHAP configuration on cluster 1:

```
VPlexcli:/> iscsi chap front-end set-default-credential -c cluster-1/ -u user
-s secret123456 --secured
Successfully set default credential on 8 ports.
```

See also

- iscsi chap front-end add-credentials
- iscsi chap front-end disable
- iscsi chap front-end enable
- iscsi chap front-end list-credentials
- iscsi chap front-end remove-credentials
- iscsi chap front-end remove-default-credential

iscsi check-febe-connectivity

Checks the front-end and back-end configurations for iSCSI connectivity.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```
iscsi check-febe-connectivity
[-c|--cluster] cluster-context
[-h|--help]
[--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Specifies the context path of the cluster at which connectivity should be checked.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

The command checks whether the front-end and back-end configurations for iSCSI connectivity are valid and meet best practices for VPLEX iSCSI configuration.

Note: This command is valid only on systems that support iSCSI devices.

Example

Shows port groups with overlapping subnets:

```
VPlexcli:/> iscsi check-febe-connectivity -c cluster-1
The following port groups have overlapping subnets:
  iscsi-port-group-2 <==> iscsi-port-group-1
  iscsi-port-group-1 <==> iscsi-port-group-2
Validate febe connectivity found errors in the system's configuration.
```

Shows an invalid sendtarget:

```
VPlexcli:/> iscsi check-febe-connectivity -c cluster-1
The following IP Addresses are used in multiple conflicting locations:
  10.10.100.150:
    director-1 BE01
    sendtarget
Validate febe connectivity found errors in the system's configuration.
```

Shows the front-end and back-end connectivity as valid:

```
VPlexcli:/> iscsi check-febe-connectivity -c cluster-1
Validate febe connectivity found no errors in the system's configuration.
```

See also


- `iscsi sendtargets add`
- `iscsi sendtargets list`
- `iscsi sendtargets rediscover`
- `iscsi sendtargets remove`

iscsi isns add

Adds iSCSI iSNS servers to the cluster, so that the cluster can register the iSCSI front-ends at those iSNS servers, allowing hosts to discover the VPLEX by querying the iSNS server.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi isns add
[-s|--sockets] server-sockets [, server-sockets,...]
[-c|--cluster] cluster-context
[-h|--help]
```

Arguments


Required arguments	
<code>[-s --sockets] <i>server-sockets</i></code>	*Specifies the IP address(es), with optional :port, of the iSNS server(s). If unspecified, the port defaults to 3205.

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of the cluster on which the iSNS server should be added.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

Adds one or more server addresses to the list of iSNS servers.

 **Note:** This command is only valid on systems that support iSCSI devices.

Example

Adding two iSNS servers with IP addresses of 192.168.100.2 and 192.168.101.2:

```
VPlexcli:/> iscsi isns add -c cluster-1 -s 192.168.100.2,192.168.101.2
Successfully added 2 iSNS servers.
```

Add two iSNS server addresses with specified ports to the list of iSNS servers.

```
VPlexcli:/> iscsi isns add -c cluster-1 -s
192.168.100.2:3000,192.168.101.2:3001
Successfully added 2 iSNS servers.
```

Add one iSNS server with IP address 192.168.100.3:

```
VPlexcli:/clusters/cluster-1> iscsi isns add -s 192.168.100.3
Successfully added 1 iSNS servers.
```

Attempt to add two iSNS servers that are already configured with the given IP addresses.

```
VPlexcli:/> iscsi isns add -c cluster-1 -s 192.168.100.2,192.168.101.2
Unable to configure the following iSNS servers:
iSNS Server      Error
-----
-----
192.168.100.2 iSNS Server 192.168.100.2:3260 is already configured at
cluster-1
192.168.101.2 iSNS Server 192.168.101.2:3260 is already configured at
cluster-1
Successfully added 0 iSNS servers.
```

See also

- `iscsi isns list`
- `iscsi isns remove`

iscsi isns list

Lists the iSNS servers configured on the cluster.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```
iscsi isns list
[-c|--cluster] cluster-context
[-h|--help]
```

Arguments

Optional Arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of the cluster for which existing iSNS servers should be listed.
<code>[-h --help]</code>	Displays command line help.

Description

Displays a list of configured iSNS servers on the specified cluster.

Note: This command is only valid on systems that support iSCSI devices. If a server is not configured on all directors, a warning message indicates the servers are missing from specific directors.

Example

List the iSNS servers on cluster-1.

```
VPlexcli:/> iscsi isns list -c cluster-1
iSNS Servers for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
  10.10.101.2  3260
```

List the iSNS servers from cluster context.

```
VPlexcli:/clusters/cluster-1> iscsi isns list
iSNS Servers for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
  10.10.101.2  3260
```

List missing iSNS servers on director 2.

```
VPlexcli:/> iscsi isns list -c cluster-1
iSNS Servers for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
```

```

10.10.101.2 3260
Warning: Some directors are missing iSNS Servers:
Director   Address   Port
-----
director-2 10.10.101.4 3260
director-2 10.10.101.2 3260
    
```

See also


- `iscsi isns add`
- `iscsi isns remove`

iscsi isns remove

Removes iSNS servers from the list of configured iSNS servers on the cluster.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```

iscsi isns remove
[-s|--sockets] server-sockets [, server-sockets,...]
[-c|--cluster] cluster-context
[-h|--help]
    
```


Arguments

Required arguments	
<code>[-s --sockets] <i>server-sockets</i></code>	*Specifies the IP address(es), with optional :port, of the iSNS servers. If unspecified, the port defaults to 3205.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of the cluster from which the iSNS server should be removed.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

Removes one or more server addresses from the list of iSNS servers.

 **Note:** This command is only valid on systems that support iSCSI devices.

Example

Remove two iSNS servers with IP addresses of 192.168.100.2 and 192.168.101.2:

```

VPlexcli:/> iscsi isns remove -c cluster-1 -s 192.168.100.2,192.168.101.2
Successfully removed 2 iSNS servers.
    
```

Remove two iSNS servers with IP addresses and specified ports.

```
VPlexcli:/> iscsi isns remove -c cluster-1 -s
192.168.100.2:3000,192.168.101.2:3001
Successfully removed 2 iSNS servers.
```

Remove one iSNS server with IP address 192.168.100.3:

```
VPlexcli:/clusters/cluster-1> iscsi isns remove -s 192.168.100.3
Successfully removed 1 iSNS servers.
```

Attempt to remove two iSNS servers that have already been removed.

```
VPlexcli:/> iscsi isns remove -c cluster-1 -s 192.168.100.2,192.168.101.2
Unable to remove the following iSNS servers:
iSNS Server      Error
-----
-----
192.168.100.2 iSNS Server 192.168.100.2:3260 is not present at cluster-1
192.168.101.2 iSNS Server 192.168.101.2:3260 is not present at cluster-1
Successfully removed 0 iSNS servers.
```

See also


- `iscsi isns add`
- `iscsi isns list`

iscsi sendtargets add

Adds iSCSI storage array target portals to specified cluster, enabling ports to be discovered.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi sendtargets add
[-s|--sockets] sendtarget-sockets [, sendtarget-sockets,...]
[-c|--cluster] cluster-context

[-h|--help]
```

Arguments


Required arguments	
<code>[-s --sockets] <i>sendtarget-sockets</i></code>	* Specifies the IP address(es), with optional <code>:port</code> , of the iSCSI sendtargets. If unspecified, the port defaults to 3260.
Optional arguments	

<code>[-c --cluster] <i>cluster-context</i></code>	Context path of the cluster on which target portals should be added.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

After an iSCSI target portal is added, discovery takes place immediately. Thereafter, periodic discoveries are automatically performed at 30-minute intervals.

 **Note:** This command is only valid on systems that support iSCSI devices.

Example

Add two target portals with IP addresses of 192.168.100.2 and 192.168.101.2:

```
VPlexcli:/> iscsi sendtargets add -c cluster-1 -t 192.168.100.2, 192.168.101.2
Successfully added 2 target portals.
```

Add two target portals with IP addresses and specified ports, 192.168.100.2:3000 and 192.168.101.2:3001.

```
VPlexcli:/> iscsi sendtargets add -c cluster-1 -t 192.168.100.2:3000,
192.168.101.2:3001
Successfully added 2 target portals.
```

Attempt to add two target portals that are already configured with the given IP addresses.

```
VPlexcli:/> iscsi sendtargets add -c cluster-1 -t 192.168.100.2, 192.168.101.2
Unable to configure the following target portals:
Target Portal Error
-----
192.168.100.2 Target portal 192.168.100.2:3260 is already configured at
cluster-1
192.168.101.2 Target portal 192.168.101.2:3260 is already configured at
cluster-1
Successfully added 0 target portals.
```

See also

- `iscsi check-febe-connectivity`
- `iscsi sendtargets list`
- `iscsi sendtargets rediscover`
- `iscsi sendtargets remove`

iscsi sendtargets list

Lists iSCSI storage array target portals in cluster.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```
iscsi sendtargets list
[-c|--cluster] cluster-context
[-h|--help]
[--force]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of cluster for which existing target portals should be listed. Defaults to local cluster if not specified.
<code>[-h --help]</code>	Displays command line help.

Description

Lists the iSCSI target portals available on at least one director. If a portal is not available on all directors, a warning message indicates which portals are missing from which directors.

Note: This command is only valid on systems that support iSCSI devices.

Example

List target portals with `-c` option:

```
VPlexcli:/> iscsi sendtargets list -c cluster-1
SendTarget Addresses for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
  10.10.101.2  3260
```

List target portals from cluster context:

```
VPlexcli:/clusters/cluster-1> iscsi sendtargets list
SendTarget Addresses for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
```

List target portals and encounter directors missing target ports:

```
VPlexcli:/> iscsi sendtargets list -c cluster-1
SendTarget Addresses for cluster-1:
  Address      Port
  -----
  10.10.100.4  3260
  10.10.100.2  3260
  10.10.101.4  3260
  10.10.101.2  3260
Warning: Some directors are missing targets:
```

```

Director   Address      Port
-----
director-2 10.10.101.4 3260
director-2 10.10.101.2 3260
    
```

See also


- `iscsi check-febe-connectivity`
- `iscsi sendtargets add`
- `iscsi sendtargets rediscover`
- `iscsi sendtargets remove`

iscsi sendtargets rediscover

Issues a rediscovery on all sendtargets on the cluster.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```


iscsi sendtargets rediscover
[-c|--cluster] cluster-context
[-h|--help]
    
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of cluster for which existing target portals should be rediscovered. Defaults to local cluster if not specified.
<code>[-h --help]</code>	Displays command line help.

Description

Issues a rediscovery on all sendtargets on the cluster.

 **Note:** This command is only valid on systems that support iSCSI devices.

Example

Rediscover all sendtargets on cluster 1:

```

Vplexcli:/> iscsi sendtargets rediscover --cluster cluster-1/
Issued sendtarget rediscovery.
    
```

See also

- `iscsi check-febe-connectivity`


- `iscsi sendtargets add`
- `iscsi sendtargets list`
- `iscsi sendtargets remove`

iscsi sendtargets remove

Removes an iSCSI storage array's target portals from the cluster list of configured target portals.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi sendtargets remove
```

```
[-s|--sockets] sendtarget-sockets [, sendtarget-sockets ,...]
```

```
[-c|--cluster] cluster-context
```

```
[-h] --help
```


Arguments

Required arguments	
<code>[-s --sockets] <i>sendtarget-sockets</i></code>	* Specifies the IP address(es), with optional :port, of the iSCSI sendtargets. If unspecified, the port defaults to 3260.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of the cluster from which the target portals should be removed.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

Displays the list of iSCSI targets that are discovered on the cluster.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

Remove two target portals with IP addresses of 192.168.100.2 and 192.168.101.2:

```
VPlexcli:/> iscsi sendtargets remove -c cluster-1 -t 192.168.100.2,
192.168.101.2
Successfully removed 2 target portals.
```

Remove two target portals with IP addresses and specified ports, 192.168.100.2:3000 and 192.168.101.2:3001.

```
VPlexcli:/> iscsi sendtargets remove -c cluster-1 -t 192.168.100.2:3000,
192.168.101.2:3001
Successfully removed 2 target portals.
```

Attempt to remove two target portals that are not configured with the given IP addresses.

```
VPlexcli:/> iscsi sendtargets remove -c cluster-1 -t 192.168.100.2,
192.168.101.2
Unable to remove the following target portals:
Target Portal Error
-----
-----
192.168.100.2 Target portal 192.168.100.2:3260 is not present at cluster-1
192.168.101.2 Target portal 192.168.101.2:3260 is not present at cluster-1
Successfully removed 0 target portals.
```

See also


- iscsi check-febe-connectivity
- iscsi sendtargets add
- iscsi sendtargets list
- iscsi sendtargets rediscover

iscsi targets list

Displays the list of iSCSI targets that are discovered on the cluster.

Contexts

All contexts.

 **Note:** This command is not supported in VPLEX.

Syntax

```
iscsi targets list
[-c|--cluster] cluster-context
```


```
[-h|--help]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of cluster for which existing targets should be listed. Defaults to local cluster if not specified.
<code>[-h --help]</code>	Displays command line help.

Description

Displays the list of iSCSI targets that are discovered on the cluster.

 **Note:** This command is valid only on systems that support iSCSI devices.

Example

List iSCSI targets for cluster 1:

```
VPlexcli:/> iscsi targets list -c cluster-1
```



```

iscsi Targets for cluster-1:
-----
Target                               SendTargets                               Login    Discovery
-----
iqn.1992-04.com.emc:cx.apm00114103168.a8 []    Logged-in Not discovered
iqn.1992-04.com.emc:cx.apm00114103168.b8 []    Logged-in Not discovered
iqn.1992-04.com.emc:cx.apm00114103169.a8 [10.10.101.4:3260, 10.10.100.4:3260] Logged-in Discovered
iqn.1992-04.com.emc:cx.apm00114103169.b8 [10.10.101.4:3260, 10.10.100.4:3260] Logged-in Discovered
  
```

The following directors are not discovered at all targets:

```

-----
director Target
-----
director-1 iqn.1992-04.com.emc:cx.apm00114103168.b8
director-2 iqn.1992-04.com.emc:cx.apm00114103168.a8
director-2 iqn.1992-04.com.emc:cx.apm00114103168.b8
director-3 iqn.1992-04.com.emc:cx.apm00114103168.a8
director-3 iqn.1992-04.com.emc:cx.apm00114103168.b8
director-4 iqn.1992-04.com.emc:cx.apm00114103168.a8
director-4 iqn.1992-04.com.emc:cx.apm00114103168.b8
  
```

See also

- `iscsi targets logout`

iscsi targets logout

Logs out of iSCSI storage array targets.

Contexts

All contexts.

Note: This command is not supported in VPLEX.

Syntax

```

iscsi targets logout
[-t|--targets] target [, target...]
[-c|--cluster] cluster-context
[-h|--help]
  
```

Arguments

Required arguments	
<code>[-t --targets] <i>target</i> [, <i>target...</i>]</code>	* Specifies the IQNs of the iSCSI targets.
Optional arguments	
<code>[-c --cluster] <i>cluster-context</i></code>	Context path of cluster from which the target portals should be logged out.
<code>[-h --help]</code>	Displays command line help.

* - argument is positional.

Description

Logs out of iSCSI targets.

Note: This command is valid only on systems that support iSCSI devices.

Example

Log out of a specified target on cluster 1:

```
VPlexcli:/> iscsi targets logout -t iqn.1992-04.com.emc.cx.amp00114103168.a8 -
c cluster-1/
Successfully logged-out of 1 targets.
```

Logout failed due to an invalid target:

```
VPlexcli:/> iscsi targets logout -t not-a-real-target -c cluster-1/
Unable to logout from the following targets:
Targets          Error
-----
not-a-real-target Firmware command error.
Successfully logged-out of 0 targets.
```

See also

- `iscsi targets list`

license install

Installs license to use the product feature.

Contexts

All contexts.

Syntax

```
license install
[-l|--license-file] license-file
[--delete-file]
[-h|--help]
[--verbose]
```

Arguments

Required Arguments	
<code>[-l --license-file]</code> <i>license-file</i>	Specifies the path to the license file, including the license file name to be validated and installed.
Optional Arguments	
<code>[--delete-file]</code>	Deletes the license file at the specified path after successful installation.
<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

The `license install` command validates and installs the license specified by the given license file. The license file must have been previously copied to the management server on which this command is run.

When the license is installed, the message `License installed.` is displayed.

Note: The license file is valid for only a specific cluster. The `VPLEXID` attribute value of `VENDOR_STRING` of the licensed feature in the license file must match the cluster TLA of the cluster on which the license file is valid. If the `license install` command is executed with a license file containing an invalid `VPLEXID` attribute value, an error is displayed.

See Also

- `license reset`
- `license show`

license reset

Removes any installed license for the product.

Contexts

All contexts.

Syntax

```
license reset
[-f|--force]
[-h|--help]
[--verbose]
```

Arguments

Optional Arguments	
<code>[-f --force]</code>	Remove installed license information without confirmation.
<code>[-h --help]</code>	Displays the usage for this command.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Note: Before executing `license reset`, confirm with customer support to ensure resetting is necessary.

The `license reset` command removes license for all installed product features on the local cluster.

If you do not specify the `-f` option, the license will be removed after a command line prompt, such as `Are you sure you wish to remove the installed product license? (Yes/No)`.

When the license is removed, the message `License removed.` is displayed.

See Also

- `license install`
- `license show`

license show

Displays details about the installed license or license specified by the given license file.

Contexts

All contexts.

Syntax

```
license show
[-l|--license-file] license-file
[-r|--raw]
[-h|--help]
[--verbose]
```


Arguments

Optional Arguments	
<code>[-l --license-file] <i>license-file</i></code>	Specifies the path to the license file whose details are to be shown.
<code>[-r --raw]</code>	Displays the raw (unformatted) license contents.
<code>[-h --help]</code>	Displays the usage for this command.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

The `license show` command displays license features that are installed on the local cluster. For example, an unlicensed product would return a message similar to `License is not installed`. A licensed product returns output containing information about various features, if installed (for example, migration and capacity).

This command also validates and displays usage of licensed features. In addition, the command displays warnings if usage exceeds the license capacity. The following attributes provide more details about the feature usage quantity.

 **Note:** Usage Intelligence is displayed only for the `VPLEX_LOCAL_CAPACITY` and the `VPLEX_METRO_CAPACITY` licenses.

- `Usage` : Displays number part of feature usage quantity.
- `Usage Unit` : Displays unit part of feature usage quantity.

For example, you may see warnings indicating that usage of a particular feature has exceeded its licensed capacity, as in the following output:

```
WARNING:

The local capacity usage '6TB' has exceeded the licensed capacity '5TB'
The metro capacity usage '2TB' has exceeded the licensed metro capacity '1TB'
```

See Also

- license install
- license reset

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] <i>name</i>	<p>* 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.</p> <p> Note: If this device will have another device attached (using the <code>device attach-mirror</code> 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.</p>
[-g --geometry] {raid-0 raid-1 raid-c}	<p>* Geometry for the new device. Valid values are <code>raid-0</code>, <code>raid-1</code>, or <code>raid-c</code>.</p> <p> CAUTION Use this command to create a RAID 1 device only if:</p> <ul style="list-style-type: none"> - None of the legs contains data that must be preserved - The resulting device will be initialized using tools on the host - The resulting device will be added as a mirror to another device
[-e --extents] <i>context-path,context-path...</i>	<p>* 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).</p>
Optional arguments	
[-d --stripe-depth] <i>depth</i>	<p>Required if <code>--geometry</code> is <code>raid-0</code>. Stripe depth must be:</p>

	<ul style="list-style-type: none"> • 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 <p>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.</p> <p>Concatenated RAID devices are not striped.</p>
<code>[-s --source-leg] <i>context-path</i></code>	When geometry argument is <code>raid-1</code> , picks one of the extents specified by the <code>--extents</code> argument to be used as the source data image for the new device. The command copies data from the <code>--source-leg</code> to the other legs of the new device.
<code>[-f --force]</code>	Create a RAID 1 device even if no <code>--source-leg</code> 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 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 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.

Examples

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/
/cclusters/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          -         -

```

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/storage-elements/extents/extent_Symm1852_AB0_1
VPlexcli:/clusters/cluster-2/storage-elements/extents> cd
VPlexcli:/> ll -p **/devices
/cclusters/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 above example if both the extents were thin-capable and from same storage array family, the RAID-1 would be thin-capable too. The virtual volume created on top of such a device can be thin-enabled.

- **Note:** The virtual volume must be built on top of a local RAID 0 device or a RAID 1 device. If you try to create a RAID C local-device with multiple children, or a device that incorporates multiple extents, the created local device is not thin-capable.

- The following example shows how a RAID-C device cannot be thin-capable:

```

VPlexcli:/clusters/cluster-1/storage-elements/extends> local-device
create --geometry raid-c -e
extent_TOP_101_1, extent_TOP_102_1 --name myLocalDevice

You are creating a raid-c local-device on top of 2 thin-capable extents
'extent_TOP_101_1, extent_TOP_102_1'.
The resulting local-device will not be thin-capable.

VPlexcli:/clusters/cluster-1/storage-elements/extends>
    
```

See also

- device attach-mirror
- local-device destroy
- local-device summary

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	
<code>[-d --devices] <i>context-path,context-path...</i></code>	* List of one or more device(s) to destroy.
Optional arguments	
<code>[-f --force]</code>	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`
- `local-device summary`

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

Optional arguments	
<code>[-c --clusters] cluster,cluster...</code>	Display information only for the specified clusters.

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.

Table 18 local device summary field descriptions

Field	Description
Health	
<code>devices</code>	Number of devices in the cluster.
<code>unhealthy</code>	Of the total number of devices in the cluster, the number whose health state is not "ok".
Visibility	
	Of the total number of devices in the cluster, the number with global or local visibility. <code>global</code> - 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. <code>local (default)</code> - Device is visible only to the local cluster.
Capacity	
<code>devices w/ space</code>	Of the total number of devices in the cluster, the number with available space.

Table 18 local device summary field descriptions (continued)

Field	Description
free capacity	Total free capacity on the cluster.
total capacity	Total capacity of the cluster.

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
    
```

See also

- ds summary
- export port summary
- export storage-view summary
- extent summary
- storage-volume summary
- virtual-volume provision

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

Optional arguments	
<code>[-s --source] <i>id</i></code>	ID of the source log to be filtered. Use the log source list command to display the list of source logs and their IDs.
<code>[-t --threshold]</code> <code>[< > =]0 - 7</code>	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>.
<code>[-c --component]</code> <i>name</i>	Component name to filter. Takes a regular expression as an argument. Plain strings are searched for in the component name.
<code>[-e --event-num] <i>id</i></code>	Used in conjunction with a specified component. An event ID to filter.
<code>[-m --message] <i>text</i></code>	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.
<code>[-n --no-consume]</code>	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 received events 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
- log filter list

log filter destroy

Removes a firmware log filter.

Contexts

All contexts.

Syntax

```
log filter destroy
[-f|--filter] filter
```

Arguments

Required arguments	
<code>[-f --filter] <i>filter</i></code>	ID of filter to delete.

Description

The filter is removed from the filter stack.

Use the `log filter list` command to display the filters configured on the system, and associated IDs of those filters.

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
- log filter list

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
- log filter destroy

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	
<code>[-s --source] <i>host:port</i></code>	* 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.
<code>[-p --password] <i>password</i></code>	
Optional arguments	The password to use for authenticating to the source.
<code>[-f --failover-source] <i>host:port</i></code>	IP address and port of the failover source to be added.

* argument is positional.

Description

 **CAUTION** For use by Dell EMC personnel only.

Creates a source for writing entries to the firmware log.

Example

```
VPlexcli:/> log source create --source 128.221.252.69:5988
Enter the source connection password:
VPlexcli:/> log source list
1. /var/log/VPlex/cli/events.log
.
.
.
6. [128.221.252.69:5988]/cpu0/log
7. [128.221.252.69:5988]/xmmg/log
```

See also

- log source destroy
- log source list

log source destroy

Destroys the specified log source.

Contexts

All contexts.

Syntax

```
log source destroy
[-s|--source] host:port
```

Arguments

Required arguments	
<code>[-s --source]</code> <i>host:port</i>	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 Dell EMC personnel only.

Example

```
VPlexcli:/> log source list
1. /var/log/VPlex/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
VPlexcli:/> log source destroy --source 128.221.252.69:5988
```

See also

- log source create
- log source list

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
- log source create

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

Optional arguments	
<code>[-v --logging-volume] <i>logging-volume</i></code>	Logging volume to which to add the mirror.
<code>[-m --mirror] {<i>name context-path</i>}</code>	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.

See also

- logging-volume create
- logging-volume destroy

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	
<code>[-n --name] <i>name</i></code>	* Name for the new logging volume.
<code>[-g --geometry] {raid-0 raid-1}</code>	* Geometry for the new volume.
<code>[-e --extents] <i>context-path,context-path...</i></code>	* 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	
<code>[-d --stripe-depth] <i>depth</i></code>	<p>Required if <code>--geometry</code> is <code>raid-0</code>. Stripe depth must be:</p> <ul style="list-style-type: none"> • 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 <p>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.</p> <p>Best practice regarding stripe depth is to follow the best practice of the underlying array.</p> <p>Concatenated RAID devices are not striped.</p>

* - 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 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.

Table 19 logging volume display fields

Field	Description
<code>application-consistent</code>	Whether or not this storage volume is application-consistent.
<code>biggest-free-segment-block-count</code>	The block count of the largest remaining free segment in the logging volume. This is the upper limit on the size of a new allocated segment.
<code>block-count</code>	The number of blocks in the volume.
<code>block size</code>	The size of a single block, in kilobytes.
<code>capacity</code>	The total number of bytes in the volume. Equals the <code>block-size</code> multiplied by the <code>block-count</code> .
<code>component-count</code>	The number of mirrors in this raid-1 logging volume.
<code>free-capacity</code>	The number of free slots for storage-volume headers in this logging volume.
<code>geometry</code>	Indicates the geometry or redundancy of this device. Will always be raid-1.
<code>health-indications</code>	If health-state is not "ok", additional information.
<code>health-state</code>	<ul style="list-style-type: none"> • <code>ok</code> - The storage volume is functioning normally. • <code>degraded</code> - 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.)

Table 19 logging volume display fields (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>unknown</code> - VPLEX cannot determine the storage volume's Health state, or the state is invalid. • <code>non-recoverable error</code> - 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. • <code>critical failure</code> - VPLEX has marked the storage volume as hardware-dead.
<code>locality</code>	<p>Locality of the supporting device.</p> <ul style="list-style-type: none"> • <code>local</code> - The volume is local to the enclosing cluster. • <code>remote</code> - The volume is made available by a different cluster than the enclosing cluster, and is accessed remotely. • <code>distributed</code> - The virtual volume either has, or is capable of having, legs at more than one cluster. *} }
<code>operational status</code>	<ul style="list-style-type: none"> • <code>ok</code> - The storage volume is functioning normally. • <code>degraded</code> - 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.) • <code>unknown</code> - VPLEX cannot determine the storage volume's Health state, or the state is invalid. • <code>error</code> - VPLEX has marked the storage volume as hardware-dead. • <code>starting</code> - The storage volume is not yet ready. • <code>lost-communication</code> - The storage volume is unreachable.
<code>rebuild-allowed</code>	Whether or not this device is allowed to rebuild.
<code>rebuild-eta</code>	The estimated time remaining for the current rebuild to complete.

Table 19 logging volume display fields (continued)

Field	Description
rebuild-progress	The percentage of this device that has been rebuilt.
rebuild-status	The rebuild status of this device.
rebuild-type	The rebuild type. <ul style="list-style-type: none"> full - A full copy of all the blocks. incremental - Uses a checksum differencing algorithm to transfer only those (chunks of) blocks that are different. comparison - A comparison copy. resync - A resync rewrites blocks that may have been affected by a director failure, guaranteeing that the mirror legs are identical.
stripe-depth	The depth of a stripe in bytes when geometry is raid-0.
supporting-device	The local, remote, or distributed device underlying the virtual volume.
system-id	Name assigned to the logging-volume.
transfer-size	The transfer size during rebuild in bytes. See About transfer-size in the batch-migrate start command.
volume-type	For logging volumes, this is always logging-volume.
/components context	
Name	Name of the extent.
Slot number	The slot number of the component.
Type	Indicates the type of component: a storage-volume, extent, or device.
Operational Status	The operational status for the entity. This indicates whether the entity is functioning, and if so, how well it is functioning.
Health State	Represents an overview of the health of the extent.
Capacity	Represents an overview of the capacity of the extent.
/segments context	
Name	Name of the segment.
Starting block	Always 0.

Table 19 logging volume display fields (continued)

Field	Description
Block count	Number of blocks in the segment.
Use	Indicates how the segment is used.

Example

```

VPlexcli:/clusters/cluster-1/system-volumes> logging-volume create -n
c1_log_vol -g raid-1 -e extent_1 , extent_2
VPlexcli:/clusters/cluster-1/system-volumes> cd c1_log_vol
VPlexcli:/clusters/cluster-1/system-volumes/c1_log_vol> ll /clusters/
cluster-1/system-volumes/c1_log_vol
/custers/cluster-1/system-volumes/c1_log_vol:
Attributes:
Name                               Value
-----
application-consistent             false
biggest-free-segment-block-count    2612155
block-count                         2621440
block-size                          4K
capacity                            10G
component-count                     1
free-capacity                       9.97G
geometry                            raid-0
health-indications                  []
health-state                        ok
locality                            local
operational-status                  ok
rebuild-allowed                     -
rebuild-eta                         -
rebuild-progress                    -
rebuild-status                      -
rebuild-type                        -
stripe-depth                        4K
supporting-device                   logging_c1_log
system-id                           logging_c1_log
transfer-size                       -
volume-type                         logging-volume
Contexts:
Name      Description
-----
components The list of components that support this logging-volume.
segments  Shows what parts of the logging volume are assigned to log changes
on distributed-device legs.
VPlexcli:/clusters/cluster-1/system-volumes/c1_log_vol> ll /clusters/
cluster-1/system-volumes/c1_log_vol/components
/custers/cluster-1/system-volumes/c1_log_vol/components:
Name      Slot      Type      Operational      Health
Capacity
-----
Number    Status      State
-----
extent_VNX-1912_LUN10_1  0      extent  ok              ok              15G
extent_VNX-1912_LUN11_1  1      extent  ok              ok              15G
VPlexcli:/clusters/cluster-1/system-volumes/c1_log_vol> ll /clusters/
cluster-1/system-volumes/c1_log_vol/segments
/custers/cluster-1/system-volumes/c1_log_vol/segments:
Name      Starting      Block
Use
-----
Block      Count

```

```

-----
-----
-----
allocated-c1_dr1ActC1_softConfig_CHM_C1_0000          1084      17
allocated for

c1_dr1ActC1_softConfig_CHM_C1_0000
allocated-c1_dr1ActC1_softConfig_CHM_C1_0001          1118      17
allocated for

c1_dr1ActC1_softConfig_CHM_C1_0001
.
.
.
allocated-r0_deviceTgt_C2_CHM_0001                    2077      17
allocated for r0_deviceTgt_C2_CHM_0001
allocated-r1_mirrorTgt_C1_CHM_0001                    2060      17
allocated for r1_mirrorTgt_C1_CHM_0001
free-1057                                              1057      10
free
free-2094                                              2094     3930066
free
free-40                                                 40         2
free
free-82                                                 82         2
free
VPLEXcli:/clusters/cluster-1/system-volumes/c1_log_vol>

```

See also

- extent create
- logging-volume add-mirror
- logging-volume destroy
- storage-volume claim

logging-volume detach-mirror

Detaches a mirror from a logging volume.

Contexts

All contexts.

Syntax

```

logging-volume detach-mirror
[-m|--mirror] mirror
[-v|--logging-volume] logging-volume
[-s|--slot] slot-number
[-h|--help]
[--verbose]

```

Arguments

Optional arguments	
--------------------	--

<code>[-m --mirror] <i>mirror</i></code>	* Specifies the name or context path of the logging volume mirror to detach. If you specify the mirror, do not specify the slot number.
<code>[-v --logging-volume] <i>logging-volume</i></code>	Specifies the name of the logging volume from which to detach the mirror.
<code>[-s --slot] <i>slot-number</i></code>	Specifies the slot number of the mirror to detach. If you specify the slot number, do not specify the mirror.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

This command detaches a mirror from a logging-volume. The logging-volume must have a RAID1 geometry and the mirror must be a direct child of the logging-volume.

You must specify the `--slot` or `--mirror` option but not both.

To detach a mirror from a component of a logging-volume use the `device detach-mirror` command.

Example

Lists the attributes of the logging volume:

```

Vplexcli:/clusters/Hopkinton/system-volumes> ll logging_vol/
/clusters/Hopkinton/system-volumes/logging_vol:
Attributes:
Name                               Value
-----
application-consistent             false
biggest-free-segment-block-count   2620324
block-count                         2621440
block-size                          4K
capacity                            10G
component-count                     1
free-capacity                       10G
geometry                            raid-1
health-indications                  []
health-state                        ok
locality                            local
operational-status                  ok
provision-type                      legacy
rebuild-allowed                     true
rebuild-eta                          -
rebuild-progress                    -
rebuild-status                      done
rebuild-type                        full
stripe-depth                        -
supporting-device                   logging
system-id                           logging
transfer-size                       128K
volume-type                         logging-volume
Contexts:
Name                               Description
-----
components  The list of components that support this logging-volume.
segments    Shows what parts of the logging volume are assigned to log
              changes on distributed-device legs.
    
```

Lists the logging volume components:

```
VPlexcli:/clusters/Hopkinton/system-volumes> ll logging_vol/components/
/clusters/Hopkinton/system-volumes/logging_vol/components:
Name                               Slot      Type      Operational  Health  Capacity
-----                               Number    -----  Status      State   -----
extent_CLARiION1389_LUN_00023_1    0         extent   ok           ok      10G
```

Specify the name of the mirror to detach:

```
VPlexcli:/clusters/Hopkinton/system-volumes> logging-volume detach-mirror --
mirror extent_CLARiION1389_LUN_00023_1/*
```

Specify the slot number of the mirror to detach:

```
VPlexcli:/clusters/Hopkinton/system-volumes> logging-volume detach-mirror --
logging-volume logging_vol/ --slot 0
```

See also

- logging-volume add-mirror
- logging-volume create
- logging-volume destroy

logging-volume destroy

Destroys an existing logging volume.

Contexts

All contexts.

Syntax

```
logging-volume destroy
[-v|--logging-volume] logging-volume
```

Arguments

Required arguments	
<code>[-v --logging-volume] <i>logging-volume</i></code>	* 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
- logging-volume create
- logging-volume detach-mirror

logical-unit forget

Forgets the specified logical units (LUNs).

Contexts

All contexts.

Syntax

```
logical-unit forget
[-s|--forget-storage-volumes]
[-u|--logical-units] context-path,context-path,...
```

Arguments

Required arguments	
<code>[-u --logical-units] <i>context-path</i></code>	List of one or more LUNs to forget.
Optional arguments	
<code>[-s --forget-storage-volumes]</code>	If a LUN has an associated storage-volume, forget it AND the associated storage-volume.

Description

Forget one or more logical units (LUNs). Optionally, forget the storage volume if one is configured on the LUN. This command attempts to forget each LUN in the list specified, logging/displaying errors as it goes.

A logical unit can only be forgotten if it has no active paths. LUNs can be remembered even if a cluster is not currently in contact with them. This command tells the cluster that the specified LUNs are not coming back and therefore it is safe to forget about them.

If a specified LUN has an associated storage-volume, that LUN is skipped (is not forgotten).

Use the `--verbose` argument to print a message for each volume that could not be forgotten.

Use the `--forget-storage-volume` argument to forget the logical unit AND its associated storage-volume. This is equivalent to using the storage-volume forget command on those storage-volumes.

Example

Forget the logical units in the current logical unit context:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-arrays/EMC-
SYMMETRIX-192602773/logical-units> logical-unit forget
13 logical-units were forgotten.
102 logical-units have associated storage-volumes and were not forgotten
```


Use the `--verbose` arguments to display detailed information about any logical units that could not be forgotten:

```

VPlexcli:/clusters/cluster-1/storage-elements/storage-arrays/EMC-
SYMMETRIX-192602773/logical-units> logical-unit forget --forget-storage-
volumes --verbose
WARNING: Error forgetting logical-unit: Logical-unit
'VPD83T3:6006016030802100e405a642ed16e111'has active paths and cannot be
forgotten.
.
.
.
WARNING: Error forgetting storage-volume
'VPD83T3:60000970000192602773533030353933': The 'use' property of storage-
volume VPD83T3:60000970000192602773533030353933' is 'meta-data' but must be
'unclaimed' or 'unusable' before it can be forgotten.
.
.
.
13 logical-units were forgotten:
VPD83T3:60000970000192602773533030353777
.
.
.
11 storage-volumes were forgotten:
VPD83T3:6006016030802100e405a642ed16e1099
.
.
.

```

See also

- `storage-volume forget`

ls

Displays information about the current object or context.

Contexts

All contexts.

Syntax

```

ls
[-l|--long]
[-a|--attributes]
[-A|--no-attributes]
[-t|--attribute] selector
[-p|--paged]
[-m|--commands]
[-f|--full]
[-C|--no-contexts]
[-x |--cache-max-age]
context, [[context]]...

```

Arguments

Optional arguments	
<code>[-l --long]</code>	Display more detailed information.
<code>[-a --attributes]</code>	Includes the attributes of the target contexts.

<code>[-A --no-attributes]</code>	Excludes attributes.
<code>[-t --attribute-selector]</code> <i>attribute-selector</i>	Display the contents of the specified attributes.
<code>[-p --paged]</code>	Page output if it is longer than the window height.
<code>[-m --commands]</code> <i>depth</i>	Includes commands in the listing.
<code>[-f --full]</code>	Do not summarize long attribute listings.
<code>[-C --no-contexts]</code>	Excludes contexts from the listing.
<code>[-x --cache-max-age]</code>	Maximum age of cached context data to be considered fresh enough. Default: 0, which always triggers a refresh.

Description

The contents of a context include: its child contexts; its attributes; and the available commands.

The context name can be any valid glob pattern.

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
default-write-same-16-template      true
default-xcopy-template              true
director-names                      [director-1-1-A, director-1-1-B]
island-id                            1
top-level-assembly                  FNM00151000986
operational-status                  ok
transition-indications              []
transition-progress                  []
health-state                         ok
health-indications                   []

Contexts:
connectivity      consistency-groups      devices      exports
performance-policies  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`

management-server set-ip

Assigns IP address, net-mask, and gateway IP address to the specified management port. Assign IP from peer MMCS.

Contexts

All contexts.

Syntax

```
management-server set-ip
```

```
[-i|--ip-netmask] destination-IP-address/ mask
```

```
[-g|--gateway] IP-address
```

```
[-p|--management-port] context-path
```

```
[-r|--peer]
```

Arguments

Required arguments	
[-i --ip-netmask] <i>destination-IP-address/ mask</i>	<p>The address and subnet mask of the Ethernet port. The format of the address/subnet mask depends on the version of IP.</p> <ul style="list-style-type: none"> 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] <i>IP-address</i>	The IP address of the gateway for this management server.
[-p --management-port] <i>context-path</i>	Ethernet port for which the parameters are assigned/changed.
[-r --peer]	Invokes the <code>set-ip</code> command on the peer MMCS. This command is only available on MMCS-A. For example, use this option to set the public IP of MMCS-B from the MMCS-A CLI.

Description

The 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 server that you can connect 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
auto-negotiate      true
duplex               auto-negotiated
gateway             10.31.52.1
inet6-address        []
inet6-gateway        []
net-mask            255.255.255.224
speed               auto-negotiated
status              ok
```

⚠ 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 clusters and the Cluster Witness server.

Additional failures (for example a remote cluster failure) while VPN between the clusters and the Cluster 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 SolVe Desktop.

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:

```
VPlexcli:/> ll /management-server/ports/eth3
/management-server/ports/eth3:
Name                Value
-----
address             10.31.52.70
auto-negotiate      true
duplex               full
gateway             10.31.52.21
inet6-address        -
inet6-gateway        -
net-mask            255.255.248.0
speed               0.977GB/s
status              up
VPlexcli:/> management-server set-ip --ip-netmask 10.31.52.197/255.255.252.0
--gateway 10.31.52.1 -p eth3
```

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:

```
VPlexcli:/> ll /management-server/ports/eth3
Name                Value
-----
address             10.31.52.197
auto-negotiate      true
duplex               full
```

```

gateway          10.31.52.1
inet6-address    [3ffe:80c0:22c:803c:215:17ff:fecc:4408/64 Scope: Global,
                 3ffe:80c0:22c:803c:415:17ff:fecc:4408/64 Scope: Link]
inet6-gateway    3ffe:80c0:22c:803c::1
net-mask         255.255.252.0
speed            0.977GB/s
status           up
Vplexcli:/> management-server set-ip --ip-netmask
3ffe:80c0:22c:803c:215:17ff:fed2:fe88/64
--gateway 3ffe:80c0:22c:803c::1 -p eth3

```

Set public IP of MMCS-B from MMCS-A:

```

Vplexcli:/>management-server set-ip
--ip-netmask 10.103.101.112/255.255.248.0 -g 10.103.96.1 -p eth3 -r

```

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	
<code>[-m --manifest] <i>pathname</i></code>	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.

See also

- `manifest version`

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`

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	
<code>[-d --storage-volume] <i>context-path</i></code>	Storage-volume to attach as a mirror to the meta-volume.
<code>[-v --meta-volume] <i>context-path</i></code>	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 volumes must be:

- Not empty.
- At the implied or specified cluster.
- Unclaimed.
- 78 GB or larger.
-

Dell EMC recommends you create a mirror and a backup of the meta-volume using at least two disks from two different arrays.

 **Note:** You can attach a mirror 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

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] <i>context-path</i>	<p>* List of two or more storage volumes to use in creating the backup meta-volume. The specified storage-volumes must be:</p> <ul style="list-style-type: none"> • Not empty • At the implied or specified cluster • Unclaimed • 78 GB or larger. • <p>Type the system IDs for multiple (two or more) storage volumes, separated by commas.</p>
Optional arguments	
[-c --cluster] <i>context-path</i>	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_YYYYMMdd_HHmms

Metadata is read from the meta-volume only during the boot of each director.

Create a backup meta-volume:

- As part of an overall system health check before a major migration or update.
- If the system permanently loses access to both meta-volumes.
 -  **Note:** No modifications should be made to the system during the backup procedure. Make sure that all other users are notified.

Use the `ll` 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
- meta-volume destroy

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	
<code>[-n --name] <i>name</i></code>	* Name of the new meta-volume.
<code>[-d --storage-volume] <i>context-path</i></code>	* 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.  Note: Specify two or more storage volumes. Storage volumes should be on different arrays.
Optional arguments	
<code>[f --force]</code>	Forces the meta-volume to be created without asking for confirmation.

* - argument is positional.

Description


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.

Dell 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.

Examples

In the following example:

- The `configuration show-meta-volume-candidates` command displays possible candidates:

 **Note:** Example out put 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
-----
-----
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
l
.
.
.
VPlexcli:/> meta-volume create --name c1_meta -storage-volumes
VPD83T3:60000970000192601707533031333136,
VPD83T3:60060480000190300487533030343445

VPlexcli:/> cd /clusters/cluster-1/system-volumes

VPlexcli:/clusters/cluster-1/system-volumes> ll c1_meta

/clusters/cluster-1/system-volumes/c1_meta:
Attributes:
Name                               Value
-----
active                             true
application-consistent             false
block-count                         20971264
block-size                          4K
capacity                            80G
component-count                     2
free-slots                          27199
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                           c1_meta
transfer-size                       128K
volume-type                         meta-volume
Contexts:
Name                               Description
-----
components The list of components that support this device or system virtual
           volume.

```

```

VPlexcli:/> configuration show-meta-volume-candidates
Name                               Capacity...   Array Name
-----
VPD83T3:6006048c460345d122db7605e8b18863 45G
EMC-Celerra-APM00140825464
VPD83T3:6006048c54c01376339c6efc9542317c 20G
EMC-Celerra-APM00140825464
VPD83T3:6006048c5b272e8526c2203c6d56dc88 100G
Celerra-APM00140825464
VPD83T3:6006048c8918b7022b8ca80efbaa7fa0 45G
EMC-Celerra-APM00140825464
VPD83T3:6006048cbc54f178bb648e9d48a3cd7c 45G
EMC-Celerra-APM00140825464
VPD83T3:6006048cbf68882ef8b8031ba611ad77 20G
EMC-Celerra-APM00140825464

```

See also

- meta-volume destroy

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	
<code>[-v --meta-volume]</code> <i>context-path</i>	- * Meta-volume to destroy.
Optional arguments	
<code>[f --force]</code>	- 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.

* - argument is positional.

Description

The meta-volume cannot be destroyed if its `active` attribute is true.

Example

In the following example:

- `ll` 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 metal
/clusters/cluster-1/system-volumes/metal:
Attributes:
Name                               Value
-----
active                             false
application-consistent            false
block-count                        23592704
.
.
.
VPlexcli:/clusters/cluster-1/system-volumes> meta-volume destroy -v metal
Meta-volume 'metal' will be destroyed. Do you wish to continue? (Yes/No) y
```

See also

- `meta-volume create`

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	
<code>[-d --storage-volume] <i>context-path</i></code>	Storage volume to detach as a mirror from the meta-volume.
<code>[-v --meta-volume] <i>context-path</i></code>	* The meta-volume from which the storage-volume/mirror should be detached.
Optional arguments	
<code>[-f --force]</code>	Force the mirror to be discarded. Required when the <code>--discard</code> argument is used.
<code>[-s --slot] <i>slot-number</i></code>	The slot number of the mirror to be discarded. Applicable only when the <code>--discard</code> argument is used.
<code>[-u --detach-unreachable-mirror]</code>	Supports the discard of an unreachable mirror.
<code>--discard</code>	Discards the mirror to be detached. The data is not discarded.

* - 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
-----
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

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] <i>context-path</i>	Storage volume to move metadata to. Target volume must be: <ul style="list-style-type: none"> • Unclaimed. • Must be 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).

This command fails if the destination meta volume has a lower number of meta data slots than required to support the current configuration. This is highly likely if the target meta-volume was manually created before Release 5.1 and has 32000 slots. Confirm this by using the ll command in the system volume context. See the troubleshooting procedures for VPLEX in the SolVe Desktop for information on fixing this problem.

See also

- meta-volume create
- meta-volume destroy

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] first
```


```
[-n|--number] number
```

```
[-c|--cluster] cluster
```

```
[-m|--meta-volume] meta-volume
```

```
--style {short|long|slow}
```


Arguments

Required arguments	
<code>[-c --cluster] cluster</code>	The cluster at which to analyze the active meta-volume. This argument may be omitted if the <code>--meta-volume</code> argument is present.
<code>[-m --meta-volume] meta-volume</code>	The meta-volume to analyze. This argument may be omitted if the <code>--cluster</code> argument is present.
<code>[-l --log] log file</code>	Full path to the log file on the management server.
<code>[-f --first] first</code>	Offset of first header to analyze.
<code>[-n --number] number</code>	Number of headers to analyze.
<code>--style {short long slow}</code>	<p>The style of analysis to do. Valid values:</p> <p><code>short</code> - Requires special firmware support available only in Release 5.0 and later.</p> <p><code>long</code> - Requires special firmware support available only in Release 5.0 and later.</p> <p><code>slow</code> - Available for all Release versions. Downloads the meta-volume headers from the meta-volume legs one at a time and compares them.</p> <p> CAUTION The <code>slow</code> option may take hours to complete on a production meta-volume.</p>

Description

An active meta-volume with an inconsistent on-disk state can lead to a data unavailability (DU) during NDU.

Best practice is to upgrade immediately after passing this meta-volume consistency check.

 **Note:** If any errors are reported, do not proceed with the upgrade, and contact Dell EMC Customer Support.

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`

monitor add-console-sink

Adds a console sink to the specified performance monitor.

Contexts

All contexts.

In context, command is `add-console-sink`.

Syntax

```
monitor add-console-sink
[-o|--format] {csv|table}
[-m|--monitor] monitor-name
[--force]
```

Arguments

Required arguments	
<code>[-m --monitor] context-path</code>	* Performance monitor to which to add a console sink.
Optional arguments	
<code>[-f --force]</code>	Forces the creation of the sink, even if existing monitors are delayed in their polling.
<code>[-o --format] {csv table}</code>	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 the 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 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`
- `monitor remove-sink`
- `monitor create`

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	
<code>[-m --monitor] <i>context-path</i></code>	* Performance monitor to which to add a console sink.
<code>[-f --file] <i>filename</i></code>	* File to which to send the sink's data.
Optional arguments	
<code>[-f --force]</code>	Forces the creation of the sink, even if existing monitors are delayed in their polling.
<code>[-n --name] <i>name</i></code>	Name for the new sink. If no name is provided, the default name "file" is applied.
<code>[-o --format] {csv table}</code>	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`
- `monitor collect`
- `monitor remove-sink`
- `report create-monitors`

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	
<code>[-m --monitor] <i>context-path,context-path</i></code>	One or more performance monitors 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):
.
.
.
```

See also

- monitor create
- report poll-monitors

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,...]`
- `[-t|--targets] context-path,context-path...`
- `[-f|--force]`

Arguments

Required arguments	
<code>[-n --name] <i>monitor-name</i></code>	* Name of the monitor. The name is appended to the director on which the monitor is configured.
<code>[-s --stats] <i>stat,[stat,...]</i></code>	* One or more statistics to monitor, separated by commas. Use the monitor stat-list command to display the available statistics.
Optional arguments	
<code>[-p --period] <i>collection-period</i></code>	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.
<code>[-d --director] context-path, context-path...</code>	* List of one or more comma-separated directors for which to display statistics.
<code>[-t --targets] context-path, context-path...</code>	List of one or more comma-separated targets for which to display statistics. Applicable only to statistics that require a target.
<code>[-f --force]</code>	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.

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

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 time stamped. 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
```

Examples

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 FE-ETH06:

```
VPlexcli:/monitoring> monitor create --name FE-ETH06-stats --director /
engines/engine-1-1/directors/director-1-1-A --stats fe-prt.* --targets
engines/engine-1-1/directors/director-1-1-A/hardware/ports/ETH06
```

See also

- `monitor add-console-sink`
- `monitor-add-file-sink`
- `monitor destroy`
- `monitor stat-list`
- `report create-monitors`

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	
<code>[-m --monitor]</code> <i>monitor-name</i>	* List of one or more names of the monitors to destroy.
Optional arguments	
<code>[-f --force]</code>	Destroy monitors with enabled sinks and bypass confirmation.
<code>[-c --context-only]</code>	Removes monitor contexts from 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.

* Argument is positional

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
- report create-monitors

monitor get-stats

Get last stats from monitors

Contexts

All

Syntax

get-stats

[m | --monitors= *context paths [, context paths>...]*]

-p | --parseable

-h | --help

--verbose

Arguments

Required arguments	
-m --monitors= <i>context paths [, context paths ...]</i>	* Get the last stats from the monitors specified by the listed context paths.
-p --parseable	Output parser-friendly stats names
Optional arguments	
-h --help	Displays the usage for this command.
--verbose	Provides more output during command execution. This may not have any effect for some commands.

* argument is positional

Description

Get last stats from monitors

The default polling frequency of System Wide perpetual monitors is 5 seconds and Virtual Volume perpetual monitors is 1 minute. So your application should tune the poll frequency (calling the REST API to get the stats from VPLEX) according to the poll frequency of the monitors. If your application is polling at a higher frequency than the monitor, your application will get redundant data or data that it has already polled.

Examples

```
monitor get-stat -monitors director-1-1-
A_PERPETUAL_vplex_sys_perf_mon_v19,director-1-1-
B_PERPETUAL_vplex_sys_perf_mon_v19
```

```
monitor get-stat -monitors director-1-1-
A_VIRTUAL_VOLUMES_PERPETUAL_MONITOR,director-1-1-
B_VIRTUAL_VOLUMES_PERPETUAL_MONITOR
```

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	
<code>[-s --sinks] <i>context-path,context-path...</i></code>	* List of one or more sinks 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
```

See also

- `monitor add-console-sink`
- `monitor add-file-sink`

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


Optional arguments	
[-c --categories] <i>category, category...</i>	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.

Use the `--categories` *categories* argument to display the statistics available in the specified category.

Use the `*` wildcard to display all statistics for all categories.

 **Note:** A complete list of the command output is available in the *Dell EMC VPLEX Administration Guide*.

Examples

```
Vplexcli:/> monitor stat-list --categories
cache                ip-com-port         rp-spl-node         fc-com-port
wrt-pacing           rp-spl-vol         fe-director         director
fe-lu                be-prt              ramf                 virtual-volume
com-cluster-io      directory           fe-prt              cg
storage-volume
OR
Use TAB key after the command
Vplexcli:/> monitor stat-list
cache                ip-com-port         rp-spl-node         fc-com-port
wrt-pacing           rp-spl-vol         fe-director         director
fe-lu                be-prt              ramf                 virtual-volume
com-cluster-io      directory           fe-prt              cg
storage-volume
```

See also

- `monitor create`
- *Dell EMC VPLEX Administration Guide*

ndu pre-check

Performs a pre-NDU validation and check.

Contexts


All contexts.

Syntax

```
ndu pre-check
```

Description

The `ndu pre-check` command should be run before you run a non-disruptive upgrade on a system to upgrade GeoSynchrony. This command runs through a number of checks to see if the non-disruptive upgrade would run into any errors in upgrading GeoSynchrony.

 **CAUTION** NDU pre-checks must be run within 24 hours before starting the NDU process.

Disclaimers for multipathing in `ndu pre-check` give time for you to validate hosts.

For more detailed information about NDU pre-checks, see the upgrade procedures for VPLEX in the SolVe Desktop.

The checks performed by `ndu pre-check` are listed in the Upgrade procedure for each software release. This procedure can be found in the VPLEX procedures in the SolVe Desktop.

See also

- `ndu start`
- `ndu recover`
- `ndu status`

ndu pre-config-upgrade

Disruptively upgrades a VPLEX 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

Optional arguments	
<code>[-u --firmware] <i>firmware-tar-file</i></code>	-Full path to director firmware package on the management server.
<code>[-i --image] <i>firmware-image-file</i></code>	-Full path to director firmware image on the management server.

Description

Disruptively upgrades a VPLEX when the VPLEX is not fully installed and configured.

 **CAUTION** This command requires the VPLEX be in a pre-config state. Specifically, do not use this procedure unless NO meta-volume is configured (or discoverable).

This command is used as part of a non-disruptive upgrade procedure for installed systems that have not yet been configured. For more information, see the upgrade procedures for VPLEX in the SolVe Desktop.

See also

- `ndu start`
- `ndu recover`
- `ndu status`

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.

See the upgrade procedure or the troubleshooting procedure in the SolVe Desktop for details of the `ndu recover` command and its use.

See also

- `ndu pre-check`
- `ndu start`
- `ndu status`

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-meta-volume-redundancy-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	
<code>[-i --image] <i>firmware-image-file</i></code>	<p>* Full path to director firmware image on the management server. For example:</p> <pre>/tmp/VPlexInstallPackages/VPlex-5.0.1.00.00.06-director-field-disk-image.tar</pre>
Optional arguments	
<code>[-t --targets] <i>targets,targets,...</i></code>	List of directors to upgrade.
<code>--force</code>	Must be specified to ignore SSD firmware version checking. (To upgrade to the same or older firmware)
<code>--check-only</code>	Check which directors will have their SSD firmware upgraded, not upgrade the firmware.
<code>--dry-run</code>	Do not perform the ssd firmware upgrade but run the same procedure as an actual install (including netbooting the directors).
<code>--skip-be-switch-check</code>	Skips the NDU pre-check for unhealthy back-end switches.
<code>--skip-cluster-status-check</code>	Skip the NDU pre-check for cluster problems (missing directors, suspended exports, inter-cluster link failure).
<code>--skip-confirmations</code>	Skip any user confirmations normally required before proceeding when there are NDU pre-check warnings.
<code>--skip-distributed-device-settings-check</code>	Skips the NDU pre-check for distributed device settings (auto-resume set to true).
<code>--skip-fe-switch-check</code>	Skips the NDU pre-check for unhealthy front-end switches.
<code>--skip-group-be-checks</code>	Skip all NDU pre-checks related to back-end validation. This includes the system configuration validation and unreachable storage volumes pre-checks.
<code>--skip-group-config-checks</code>	Skip all NDU pre-checks related to system configuration. This includes the system configuration validation and director commission pre-checks.

<code>--skip-group-fe-checks</code>	Skip all NDU pre-checks related to front-end validation. This includes the unhealthy storage views and storage view configuration pre-checks.
<code>--skip-group-health-checks</code>	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.
<code>--skip-meta-volume-backup-check</code>	Skips the check to verify that backups for the meta-data volumes at all clusters have been configured.
<code>--skip-meta-volume-redundancy-check</code>	Skips the NDU pre-check for verifying the meta-volume redundancy
<code>--skip-remote-mgmt-version-check</code>	Skip the remote management server version check.
<code>--skip-storage-volumes-check</code>	Skip the NDU pre-check for unreachable storage volumes.
<code>--skip-sysconfig-check</code>	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.
<code>--skip-view-config-check</code>	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.
<code>--skip-view-health-check</code>	Skip the NDU pre-check for unhealthy storage views.
<code>--skip-virtual-volumes-check</code>	Skip the NDU pre-check for unhealthy virtual volumes.
<code>--skip-wan-com-check</code>	<p>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.</p> <p> CAUTION Skipping the WAN communications pre-check may increase the risk for NDU failure should the inter-cluster communication connection fail.</p> <p> Note: Multiple skip options can be specified to skip multiple pre-checks. Enter skip options separated by a space.</p>

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.

Refer to the upgrade procedure in the SolVe Desktop for more information on using this command.

See also

- ndu start
- The VPLEX procedures in the SolVe Desktop to upgrade/troubleshoot GeoSynchrony.

ndu start

Begins the non-disruptive upgrade (NDU) process of the director firmware.

Contexts

All contexts.


Syntax

ndu start

`[--io-fwd-ask-for-confirmation] prompt type [-u|--firmware] firmware-tar-file [optional-argument [optional-argument]]`

Arguments

Required arguments	
<code>[-u --firmware] <i>firmware-tar-file</i></code>	<p>* Full path to director firmware package on the management server.</p>
<code>[--io-fwd-ask-for-confirmation] <i>prompt type</i></code>	<p>The type of the prompt that you want to see during the IO forwarding phase of the NDU. The available options are:</p> <ul style="list-style-type: none"> • <code>always</code> - Choose this option if you have hosts that require manual scanning for the paths to be visible. Assistance from the customer is required to verify that initiator paths on the hosts are alive. If the path is unavailable, resolve the issue within the timeout period that you have specified. The prompts for this options are: <ul style="list-style-type: none"> ▪ <code>Continue</code>: NDU continues even when there are missing initiator logins. Make sure that the customer is aware that missing logins can cause DU. ▪ <code>Rollback</code>: NDU rolls back and DU is avoided. The customer can check the host, resolve the issue that led to the missing initiator logins, and rerun the NDU. ▪ <code>Refresh</code>: Get the new list of initiators. If all the initiators are logged in, VPLEX displays the prompts to move forward. • <code>on-missing-logins</code> - Assistance from the customer is required to determine whether any missing initiators are from critical hosts. If paths are unavailable from critical hosts, the customer will need to resolve the issue before continuing with the NDU. The prompts for this options are: <ul style="list-style-type: none"> ▪ <code>Continue</code>: NDU continues even when there are missing initiator logins. Make sure that the customer is aware that missing logins can cause DU. ▪ <code>Rollback</code>: NDU rolls back and DU is avoided. The customer can check the host, resolve the issue that led to the missing initiator logins, and rerun the NDU.

	<ul style="list-style-type: none"> ▪ Refresh: Get the new list of initiators. After all the initiators are logged in, NDU continues without displaying any prompt. • never - No interaction is required or the customer is not available to check the host connectivity. NDU waits for all the initiators to log back in within the specified timeout period. Resolve any issues within this period. If VPLEX identifies any missing logins after the timeout period, NDU is rolled back. It is important to check whether this value must be modified according to the environment requirements of the customers and weighed against the risks.
Optional arguments	
<code>--io-fwd-timeout= <i>time</i></code>	<p>The period after which the I/O forward phase times out. In the I/O forward phase, the I/Os that are serviced to the first set of directors are forwarded to the second set of directors. The hosts are expected to connect back to the first set of directors during this period. By default, this phase lasts for 180 minutes. You can set this timeout period to a minimum of 6 minutes and a maximum of 12 hours. Use:</p> <ul style="list-style-type: none"> • s for seconds • m for minutes • h for hours • d for days
<code>--cws-package <i>cws-firmware-tar-file</i></code>	<p>Full path to Cluster Witness Server package on the management server.</p> <p> Note: Not required if upgrading to an official product release.</p>
<code>--force</code>	Ignore manifest checking of supported upgrades.
<code>--skip-cws-upgrade</code>	Skips the upgrade of the Cluster Witness Server and proceeds with the rest of the NDU.
<code>--skip-be-switch-check</code>	Skips the NDU pre-check for unhealthy back-end switches.
<code>--skip-cluster-status-check</code>	Skips the NDU pre-check for cluster problems (missing directors, suspended exports, inter-cluster link failure, and so on).
<code>--skip-confirmations</code>	Skips any user confirmations normally required before proceeding when there are NDU pre-check warnings.
<code>--skip-distributed-device-settings-check</code>	Skips the NDU pre-check for distributed device settings (auto-resume set to true).
<code>--skip-fe-switch-check</code>	Skips the NDU pre-check for unhealthy front-end switches.
<code>--skip-group-be-checks</code>	Skips all NDU pre-checks related to back-end validation. This includes pre-checks for system configuration validation and unreachable storage volumes.
<code>--skip-group-config-checks</code>	Skips all NDU pre-checks related to system configuration. This includes the system configuration validation and director commission pre-checks.

<code>--skip-group-fe-checks</code>	Skips all NDU pre-checks related to front-end validation. This includes the unhealthy storage views and storage view configuration pre-checks.
<code>--skip-group-health-checks</code>	Skips 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.
<code>--skip-meta-volume-backup-check</code>	Skips the check to verify that backups for the meta-data volumes at all clusters have been configured.
<code>--skip-meta-volume-redundancy-check</code>	Skips the NDU pre-check for verifying the meta-volume redundancy.
<code>----skip-storage-volumes-check</code>	Skip the NDU pre-check for unreachable storage volumes.
<code>--skip-sysconfig-check</code>	Skips 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.
<code>--skip-view-config-check</code>	Skips 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.
<code>--skip-view-health-check</code>	Skips the NDU pre-check for unhealthy storage views.
<code>--skip-virtual-volumes-check</code>	Skips the NDU pre-check for unhealthy virtual volumes.
<code>--skip-wan-com-check</code>	Skips the inter-cluster communications connectivity NDU pre-check and proceed with NDU even if there are errors specifically related to inter-cluster communications connectivity.
<code>--skip-local-com-check</code>	Skips the intra-cluster communications connectivity NDU pre-check and proceed with NDU even if there are errors specifically related to intra-cluster communications.
<code>--skip-total-number-of-volumes-check</code>	Skips the NDU pre-check for total number of volumes.
<code>--skip-inter-director-mgmt-connectivity-check</code>	Skips the NDU pre-check for inter-director management connectivity.
<code>--do-not-verify-wanlink-after-upgrade</code>	NDU does not check whether first-upgraders see each other on the WAN link after the upgrade.
<code>--skip-storage-view-missing-lun0-check</code>	Skips the NDU pre-check for storage-views with virtual-volumes that do not have a LUN0 ID.
<code>skip-recoverpoint-cluster-check</code>	Skips the NDU pre-check for verifying the basic health of the RecoverPoint cluster.

Description

This command starts a non-disruptive upgrade and can skip certain checks to push a non-disruptive upgrade when the `ndu pre-checks` command fails. 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.

Note: It is recommended that you upgrade VPLEX using the upgrade procedure found in the SolVe Desktop. This procedure also details when the `ndu start` command should be used with skip options and how to select and use those skip options.

See also

- `ndu pre-check`
- `ndu recover`
- `ndu status`
- The VPLEX procedures to upgrade/troubleshoot GeoSynchrony in the SolVe Desktop .

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.

See also

- `ndu pre-check`
- `ndu start`
- `ndu recover`

- The VPLEX procedures in the SolVe Desktop to upgrade/troubleshoot GeoSynchrony.

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] file
```

```
[-f|--force]
```

Arguments

Required arguments	
<code>[-m --modified-events-file] <i>file</i></code>	Path to the file containing the modified call-home events.
Optional arguments	
<code>[-f --force]</code>	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.

Use the `set` command to enable/disable call-home notifications.

Use the `ls notifications/call-home` command to display whether call-home is enabled.

See also

- `notifications call-home remove-event-modifications`
- `notifications call-home view-event-modifications`
- `notifications call-home test`

notifications call-home remove-event-modifications

Removes customized call-home events files, including customer-specific modifications and modifications recommended by Dell 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	
<code>[-c --customer-specific]</code>	If a customer-specific call-home events file has been imported, removes the file.
<code>[-e --emc-generic]</code>	If an Dell EMC call-home events file has been imported, removes the file.
<code>[-f --force]</code>	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:

- Dell EMC-generic events are modifications recommended by Dell EMC. Dell 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. Dell EMC provides a custom events file developed by Dell EMC engineering and applied by Dell 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 `ndu upgrade-mgmt-server` command to re-import the file.

See also

- `notifications call-home view-event-modifications`
- `notifications call-home test`

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	
<code>[-c --customer-specific]</code>	Displays customer specific modifications.
<code>[-e --emc-generic]</code>	Displays Dell 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 Dell EMC-generic events files, the setting in the customer-specific file overrides the entry in the Dell 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 Dell EMC generic modified events.

Use the `--verbose` argument to display detailed information.

See also

- `notifications call-home remove-event-modifications`

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 Dell 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 accounts specified to receive notifications to verify the test event arrived.

If call-home is configured to send event notifications to Dell EMC, contact Dell EMC Support to verify that the test event arrived.

Use the `set` command to enable or disable call-home notifications.

Use the `ls notifications/call-home` command to verify that call-home is enabled.

See also

- configuration event-notices-reports config
- configuration event-notices-reports reset
- notifications snmp-trap create
- set

notifications job cancel

Cancels jobs in a Queued state from the job queue.

Contexts

All contexts.

In /notifications/jobs context, command is cancel.

Syntax

```
notifications job cancel
[-j|--jobs] job-context-path [, job-context-path...]
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-j --jobs] <i>job-context-path</i></code>	* Specifies the jobs to cancel.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

- * argument is positional.

Description

This command cancels jobs that are queued in the job queue.

Only jobs that are in a Queued state can be cancelled. Once they are cancelled, the jobs remain in the job queue with a state of Cancelled.

See also

- notifications job delete
- notifications job resubmit

notifications job delete

Permanently removes jobs from the job queue.

Contexts

All contexts.

In /notifications/jobs context, command is delete.

Syntax

```
notifications job delete
[-s|--state] job-state
[-j|--jobs] job-context-path [, job-context-path...]
[-f|--force]
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-j --jobs] <i>job-context-path</i></code>	* Specifies the jobs to delete.
Optional arguments	
<code>[-s --state] <i>job-state</i></code>	Specifies a job state by which to filter the jobs. If specified, this argument filters the jobs by their states. This option is most useful when all jobs are specified in the command invocation.
<code>[-f --force]</code>	Deletes the jobs without asking for confirmation.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution.

- * argument is positional.

Description

This command permanently removes jobs from the job queue. Jobs that are in progress cannot be deleted.

A job that is in progress cannot be deleted. All other job states - cancelled, failed, and successful - can be deleted.

Example

Delete jobs that are in a failed state:

```
VPlexcli:/notifications/jobs> delete * --state failed
WARNING: The following items will be deleted:
Context
-----
/notifications/jobs/Provision_4_08-04-14-17-34-51
/notifications/jobs/Provision_4_14-04-14-15-57-03
Do you wish to proceed? (Yes/No) y
Deleted 2 jobs.
```

Shows the attempt to delete two specific jobs failed (jobs were skipped):

```
VPlexcli:/notifications/jobs> delete Provision_4_14-04-14-15-57-03,
Provision_4_08-04-14-17-34-51
WARNING: The following items will be deleted:
Context
-----
/notifications/jobs/Provision_4_08-04-14-17-34-51
```

```
/notifications/jobs/Provision_4_14-04-14-15-57-03
Do you wish to proceed? (Yes/No) y
Skipped 2 jobs.
```

Use the `--verbose` option to display more information for a job that is skipped. In this example, one job was skipped, or not deleted, because it was in progress:

```
VPlexcli:/notifications/jobs> delete * --verbose
WARNING: The following items will be deleted:
Context
-----
/notifications/jobs/Provision_4_08-04-14-17-34-51
/notifications/jobs/Provision_4_14-04-14-15-57-03
Do you wish to proceed? (Yes/No) y
WARNING: Operation Delete failed for target 'Provision_4_08-04-14-17-34-51'.
Cannot delete job 'Provision_4_08-04-14-17-34-51' because its state is
InProgress.
Deleted 1 job.
Skipped 1 job.
```

See also

- notifications job cancel
- notifications job resubmit

notifications job resubmit

Resubmits a failed or canceled job.

Contexts

All contexts.

In `/notifications/jobs context`, command is `resubmit`.

Syntax

```
notifications job resubmit
[-j|--jobs] job-context-path [, job-context-path...]
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-j --jobs] <i>job-context-path</i> [, <i>job-context-path...</i>]</code>	* Specifies the jobs to resubmit.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

- * argument is positional.

Description

Resubmit failed or canceled jobs to the job queue. The jobs will be placed in a queued state and executed in normal priority order.

See also

- `notifications job cancel`
- `notifications job delete`

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	
<code>[-n --name] <i>trap-name</i></code>	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`

- `notifications snmp-trap destroy`
- `set`

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

Optional arguments	
<code>[-s --snmp-trap] <i>trap-name</i></code>	Name of the SNMP trap sink to destroy.
<code>[-f --force]</code>	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`

password-policy reset

Resets the password-policies to the default factory settings.

Contexts

```
/security/authentication
```

Syntax

```
password-policy reset
```

```
[-h|--help]
```

```
[-f|--force]
```

```
[--verbose]
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command.
<code>[-f --force]</code>	Forces reset of password-policy configuration without asking for confirmation.
<code>[--verbose]</code>	Provides more output during command execution.

Description

Resets the password-policies to the default factory settings.

Note: This command can only be run by the `admin` user. So it will ask for the admin password.

After successful authentication, if you do not specify the `--force` option during command execution, the command prompts with a confirmation message and, based on user input, it will proceed.

Note: Using this command will override all existing password-policy configurations and return them to the default settings.

```
Vplexcli:/> security authentication password-policy reset
admin password:
The password-policy reset command will override ALL of the existing password-
policy configuration changes made to the system.
Do not run this command unless you are certain that it is required to return
password-policy to default-state.
Do you want to proceed? (Yes/No) yes
```

See also

- password-policy set

password-policy set

Each attribute in the Password policy is configurable. The new value will be updated to the respective configuration file and existing users will be updated with this configuration.

Contexts

```
/security/authentication/password-policy
```

Syntax

```
set
[password-minimum-length
| minimum-password-age
| maximum-password-age
| password-warn-days
| password-inactive-days ] value
```


Arguments

[password-minimum-length <i>value</i>]	Minimum number of characters for a password. This policy is enforced during password changes and user creation. This value can be any value between 6 and 99999. The default value is 8.
[password-minimum-days <i>value</i>]	Minimum number of days a password can be changed since the last password change. The default value is 1. Service account default is 0 days. The value can be any value between -1 and 99999.
[maximum-password-days <i>value</i>]	Maximum number of days since the last password change a password can be used. The default for maximum-password-age is 90. Service account default is 3650 days. The value can be any value between 1 and 99999.
[password-warn-days <i>value</i>]	Number of days before the password expiry. At this time, a warning message will be displayed indicating the impending password change. The default value for password-warn-days is 15. Service account default is 30 days. Password warning days can be any value between 1 and 99999 but the number of days should not be greater than password-maximum-days.
[password-inactive-days <i>value</i>]	Number of inactive days after a password expired, before the account is locked. During the inactive days you are forced to change this password first before proceeding to access the Management Server. After the inactive days, the user account will be locked. The default value for this variable is 1. The password-inactive-days can be any value between -1 and 99999.

Description

The password policies are not applicable to users configured through an LDAP server.

Password inactive days is not applied to admin user to protect the admin user from account lockouts.

 **Note:** Password Policy can be configured only by the admin user.

Example

To view the existing password-policy settings for the admin account and for other local users , run the `ll` command in the `security/authentication/password-policy` context:

```
VPlexcli:/security/authentication/password-policy> ll
Name                                     Value
-----
password-inactive-days                   1
password-maximum-days                    90
password-minimum-days                    1
password-minimum-length                   8
password-warn-days                        15
```

To view the existing password-policy settings for the service account, run the `ll` command in the `security/authentication/password-policy/service` context:

```
VPlexcli:/security/authentication/password-policy/service>
Name                                     Value
-----
password-maximum-days                    3650
password-minimum-days                    0
password-warn-days                        30
```

To view the default password-policy settings, run the `ll` command in the `security/authentication/password-policy/default` context:

Name	Value
password-inactive-days	1
password-maximum-days	90
password-minimum-days	1
password-warn-days	15

To view the permissible values for the password-policy set command, enter the command with no options:

```

VPlexcli:/security/authentication/password-policy> set
attribute          input-description
-----
name              Read-only.

password-maximum-days  Takes an integer value between 1 and 99999.
password-minimum-days  Takes an integer value between -1 and 99999. Specifying a '-1' value
will disable the policy.
password-minimum-length Takes an integer value between 6 and 99999.
password-warn-days     Takes an integer value between -1 and 99999. Specifying a '-1' value
will disable the policy.
    
```

To set the password inactive days value to three days, use this command:

```

VPlexcli:/security/authentication/password-policy> set password-inactive-days
3
admin password:
VPlexcli:/security/authentication/password-policy> ll
Name          Value
-----
password-inactive-days      3
password-maximum-days      90
password-minimum-days      1
password-minimum-length    8
password-warn-days        15
    
```

See also

- password-policy reset

plugin addurl

Adds an URL to the plug-in search path.

Contexts

All contexts.

Syntax

```


plugin addurl
[-u|--urls] url,url...
    
```

Arguments

Required arguments	
--------------------	--

<code>[-u --urls] <i>url, url...</i></code>	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`
- `plugin register`

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-il8n-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-adapt.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-il8n-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-il8n-fr.jar, file:/opt/emc/Vplex/apache-
tomcat-6.0.x/lib/tomcat-dbc.jar
```

See also

- `plugin addurl`

- `plugin register`

plugin register

Registers a shell plugin by class name.

Contexts

All contexts.

Syntax

```
plugin register
```

```
[-c|--classes] class-name[,class-name ...]
```


Arguments

Required arguments	
<code>[-c --classes] <i>class-name</i>[, <i>class-name</i>...]</code>	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`
- `plugin listurl`

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/director-1-1-A> pushd /engines/
engine-1-1/directors/director-1-1-B
[/engines/engine-1-1/directors/director-1-1-B, /engines/engine-1-1/
directors/director-1-1-A, /clusters/cluster-1/storage-elements/storage-
arrays, /, /]
VPlexcli:/engines/engine-1-1/directors/director-1-1-B> popd
[/engines/engine-1-1/directors/director-1-1-B, /clusters/cluster-1/
storage-elements/storage-arrays, /, /]
VPlexcli:/engines/engine-1-1/directors/director-1-1-A>
```

See also

- `pushd`

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

Optional arguments	
<code>[-c --context] context</code>	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, use the `pushd` command to push the first context onto the context stack:

```
VPlexcli:/>
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/director-1-1-A/
```



```
[/engines/engine-1-1/directors/director-1-1-A, /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/director-1-1-A> pushd
[/clusters/cluster-1/storage-elements/storage-arrays, /engines/engine-1-1/
directors/director-1-1-A, /, /]
VPlexcli:/clusters/cluster-1/storage-elements/storage-arrays> pushd
[/engines/engine-1-1/directors/director-1-1-A, /clusters/cluster-1/storage-
elements/storage-arrays, /, /]
VPlexcli:/engines/engine-1-1/directors/director-1-1-A>
```

See also

- `popd`

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] <i>context-path</i>	* List of one or more devices for which to change the transfer size. Wildcards are permitted. Entries must be separated by commas.
[-l --limit] <i>limit</i>	* 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 in the batch-migrate start command.

* - argument is positional.

Description

If the target devices 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`
- `rebuild status`

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
```

Arguments

Optional arguments	
<code>[-r --devices] <i>context-path...</i></code>	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`
- `rebuild status`

rebuild status

Displays all global and cluster-local rebuilds along with their completion status.

Contexts

All contexts.

Syntax

```

rebuild status
[--show-storage-volumes]

```

Arguments

Optional arguments	
<code>--show-storage-volumes</code>	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:

```

rebuild/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
```

See also

- `rebuild show-transfer-size`

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
```

Arguments

Optional arguments	
<code>[-d --directory]</code> <i>directory</i>	Directory in which to create the .csv files. Default directory path: <code>/var/log/VPlex/cli/reports/</code> 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 typeReportMonitor_director.csv
```

All other files in the directory will be ignored.

Aggregate report filenames are in the following format:

```
report typePerformance_cluster.csv
```

If an aggregate report already exists, it will be overwritten.

Examples

```
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
- report capacity-clusters
- report capacity-hosts
- report create-monitors
- report poll-monitors

report capacity-arrays

Generates a capacity report.

Contexts

All contexts.

Syntax

```
report capacity-arrays
```

```
[-t|--tier-regex] regular-expression
```

```
[-d|--directory] directory
```

Arguments

Optional arguments	
<code>[-t -tier-regex] <i>regular-expression</i></code>	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 <code>/system-defaults::tier-regular-expression</code>
<code>[-d -directory] <i>directory</i></code>	Directory in which to create the csv files. Output is written to two files: <ul style="list-style-type: none"> File for local storage: <code>CapacityArraysLocal.csv</code>. File for shared storage: <code>CapacityArraysShared.csv</code>. Default directory path: <code>/var/log/VPlex/cli/reports/*</code> 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 files already exist, the report is appended to the end of the files.

i **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.

Example

To display the raw report file, exit to the management server:

```
VPlexcli: /> exit
Connection closed by foreign host.
service@ManagementServer: ~>
```

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
.
.
.
```

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
```

To display the file with the `report capacity-arrays` command:

```
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
  Tier
  StorageVolumes (GiB)
  no-tier
  10596.152
  cluster-2
  Unallocated
  Tier
  StorageVolumes (GiB)
  no-tier
  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
- report capacity-hosts

report capacity-clusters

Generates a capacity report for every cluster.

Contexts

All contexts.

Syntax

```
report capacity-clusters
[-d|--directory] directory
[--verbose]
```

Arguments

Optional arguments	
[-d --directory] <i>directory</i>	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.
--verbose	If the --verbose argument is used: <ul style="list-style-type: none"> • Storage_volumes and volumes are listed before the summary is printed.

- **Two additional files are created:** `CapacityClustersVolumes.csv` and `CapacityClustersStorageVolumes.csv`.

Description

The capacity report 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.

Examples

```

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`
- `report capacity-hosts`

report capacity-hosts

Generates a host capacity report.

Contexts

All contexts.

Syntax

```

report capacity-hosts
[-d|--directory] directory
[--verbose]

```

Arguments

Optional arguments	
<code>[-d --directory] <i>directory</i></code>	Directory in which to create the csv files. Output is written to a file named <code>CapacityHosts.csv</code> . Default directory path: <code>/var/log/VPlex/cli/reports/</code> on the management server.
<code>--verbose</code>	If the <code>--verbose</code> argument is used, an additional file is created: <code>CapacityHostsViews.csv</code> .

Description

The host capacity information includes:

- Number of views.
- Total exported capacity in GB.
- Number of exported virtual volumes per cluster.

Example

Generate a host capacity report.

```
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`
- `report capacity-arrays`

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]
```

Arguments

Optional arguments	
<code>[-d --directory]</code> <i>directory</i>	Directory in which to create the csv files. Default directory path: <code>/var/log/VPlex/cli/reports/</code> on the management server.
<code>--force</code>	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_nn_diskReportMonitor`
- `Cluster_n_Dir_nn_portReportMonitor`
- `Cluster_n_Dir_nn_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_nn.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

Examples

In the following example:

- The `report create-monitors` command creates a `diskReportMonitor`, `portReportMonitor`, and `volumeReportMonitor` for each director
- 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
-----  -----  -----  -----  -----
-----  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      Bucket  Bucket  Bucket  Bucket
-----  -----  -----  -----  -----
-----  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`
- `monitor create`
- `monitor destroy`
- `monitor remove-sink`
- `report poll-monitors`

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
- report create-monitors

rm

Deletes a file from the corresponding share location.

Contexts

This command can only be executed in the `in` or `out` sub-contexts within the `share` context of the management server (either `/management-server/share/in` or `/management-server/share/out`).

Syntax

```
rm -n|--filename filename [-h | --help] [--verbose]
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

The `rm` command is used to delete a file from an SCP directory.

As part of Role-based access implementation, users other than **service** are not allowed shell access and access by SCP is restricted to a single directory. The SCP directory, `/diag/share/` consists of two sub-directories `in` and `out` which contain only files that can be transferred by SCP to and from of the management-server respectively.

`mangement-server/share/in` and `mangement-server/share/out` are contexts corresponding to the `in` and `out` sub-directories of the SCP directory. Users without shell access use `ls` and `rm` commands to files transferred to and from the management server with SCP.

service and **admin** users are authorized to delete any existing file in the SCP sub-directories. Other users are only authorized to delete files to which they have access.

See also

- `user add`

rp import-certificate

Imports a RecoverPoint security certificate from the specified RPA cluster.

Contexts

All contexts.

Syntax

```
rp import-certificate
```

Arguments

None.

Description

This command runs an interview script to import the RecoverPoint security certificate.

In Metro systems, run this command on both management servers.

CAUTION This command restarts VPLEX CLI and GUI sessions on the VPLEX cluster to which the RPA cluster is attached. With this release, using this command will lead to loss of VPLEX Integrated Array Services (VIAS) provisioning jobs created.

Before you begin, you will need the IP address of the RecoverPoint cluster from which to import the security certificate.

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 ...
Stopping EMC Vplex Management Console: Connection closed by foreign host.
service@sms-advil-2:/opt/emc/Vplex/tools/utills>
```

See also

- `rp summary`
- `rp validate-configuration`

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-management-address
[-u|--admin-username] admin-username
[-c|--cluster] cluster-id
```

Arguments


Required arguments	
<code>[-o --host] <i>rpa-management-address</i></code>	* The RPA cluster management IP address.
<code>[-u --admin-username] <i>admin-user-name</i></code>	* The administrative username of the RPA
Optional arguments	
<code>[-c --cluster] <i>cluster-id</i></code>	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" in the security ipsec-configure command.

* 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. Configuration of RPAs is not permitted during VPLEX NDU.

After the RPA cluster is added, information about the RPA cluster and its consistency groups and volumes appear in the following VPLEX CLI contexts and commands:

- `/recoverpoint/rpa-clusters/ip_address/volumes`
- `/clusters/cluster_name/consistency-groups/cg_name/recoverpoint`
- `rp summary` command
- `rp validate-configuration` command

Table 20 recoverpoint display fields

Field	Description
In /recoverpoint/rpa-clusters and rpa-clusters/ip-address contexts	
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 <code>rp-health-status</code> 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. OK - All components of the RP configuration are operating as expected. error - One or more components of the RP configuration is not operating as expected. The <code>rp-health-indications</code> field displays additional information. warning - One or more components of the RP configuration is not operating as expected. The <code>rp-health-indications</code> field displays additional information. unknown - 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:

Table 20 recoverpoint display fields (continued)

Field	Description
	MajorVersion.MinorVersion.ServicePack.Patch (branch.build)
In /recoverpoint/rpa-clusters/ip-address/consistency-groups context	
copies	Contains the production copy and the replica copies of which this RecoverPoint consistency group consists.
links	Contains all the communication pipes used by RecoverPoint to replicate consistency group data between the production copy and the replica copies.
replication-sets	Contains all the replication sets of which this RecoverPoint consistency group consists.
In /recoverpoint/rpa-clusters/ip-address/volumes context	
Name	Name of the volume as it appears in VPLEX.
RPA Site	See RPA Site field description above.
RP Type	<p>Role the volume is serving in RecoverPoint replication.</p> <p>Journal - 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>Replication - All volumes of a replication set. One production volume and one or two replica volumes.</p> <p>Repository - 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 RecoverPoint 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"> After a temporary failover, the Production Source becomes the Replica at Production and the Remote Replica (or

Table 20 recoverpoint display fields (continued)

Field	Description
	<p>Local Replica) becomes the Remote (or Local) Source.</p> <ul style="list-style-type: none"> After a permanent failover: the Replica at Production becomes the Production Source and the Remote (or Local) Source becomes the Remote (or Local) Replica. <p>Local Replica - The replica at the local site. Also, the role of the local copy, before failover.</p> <p>Production Source - The role of a consistency group production copy, before failover.</p> <p>Remote Replica - 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"> Cache-mode property is synchronous Consistency groups with the “visibility” property set to both clusters must also have their “storage-at-clusters” set to both clusters. Recoverpoint-enabled property set to true.
Capacity	Capacity of the volume.
In /recoverpoint/rpa-clusters/ip-address/volumes/volume context	
rp-consistency-group	See RP Group field description above.
rp-consistency-group- copy	<p>The RecoverPoint copy to which the volume belongs.</p> <p>In CDP and CRR configurations, there is one production copy and one replica copy.</p> <p>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).</p>

Table 20 recoverpoint display fields (continued)

Field	Description
	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.
rp-replication-set	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.
rp-role	See RP Role field description above.
rp-type	See RP Type field description above.
rpa-site	See RPA Site field description above.
rpa-software-serial-d	Software serial ID for each rpa-cluster.
size	See Capacity field description above.
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.

Examples

Add a RecoverPoint RPA cluster:

```
VPlexcli:/> rp rpa-cluster add -o 10.108.65.217 -u admin -c cluster-1
Enter rpa-cluster administrative password: Admin-password
Enter rpa-cluster administrative password again for verification: Admin-
password
```

Display RPA clusters.

- `ll /recoverpoint/rpa-clusters` displays summarized information about the RPA cluster.
- `ll /recoverpoint/rpa-clusters/ip-address` displays detailed information about the RPA cluster
- `ll /recoverpoint/rpa-clusters/ip-address/consistency-groups` displays summarized information about consistency groups in the RPA cluster.

- `ll /recoverpoint/rpa-clusters/ip-address/consistency-groups/cg-name` displays detailed information about the specified consistency group.
- `ll /recoverpoint/rpa-cluster/ip-address/volumes/` displays 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
-----
10.6.210.87  cluster-1     Tylenol-1  RPA 1   4.1(d.147)
10.6.211.3   cluster-2     Tylenol-2  RPA 1   4.1(d.147)
Vplexcli:/> ll recoverpoint/rpa-clusters/10.6.210.87/
/recoverpoint/rpa-clusters/10.6.210.87:
Attributes:
Name          Value
-----
admin-username      admin
config-changes-allowed true
rp-health-indications [Problem detected with RecoverPoint RPAs and splitters]
rp-health-status    warning
rp-software-serial-id -
rpa-host            10.6.210.87
rpa-id              RPA 1
rpa-site            Tylenol-1
rpa-version         4.1(d.147)
vplex-cluster       cluster-1
Contexts:
Name          Description
-----
consistency-groups Contains all the RecoverPoint consistency groups which consist of copies local to this VPLEX cluster.
volumes local Contains all the distributed virtual volumes with a local extent and the local virtual volumes which are used by this RPA cluster for RecoverPoint repository and journal volumes and replication volumes.
Vplexcli:/> cd recoverpoint/rpa-clusters/10.6.210.87/consistency-groups
Vplexcli:/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups> ll
Name
----
CG-1
CG-2
CG-3
Vplexcli:/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups> ll
CG-1
/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1:
Attributes:
Name          Value
-----
active-replicating-rp-cluster Tylenol-1
distributed-group false
enabled true
preferred-cluster Tylenol-1
preferred-primary-rpa RPA2
production-copy Pro-1
protection-type MetroPoint Replication
uid 7a59f870
Contexts:
Name          Description
-----

```

```

-----
copies          Contains the production copy and the replica copies of
which

links          Contains all the communication pipes used by
RecoverPoint to
copy           replicate consistency group data between the production
and the replica copies.
replication-sets Contains all the replication sets of which this
RecoverPoint consistency group consists.
Vplexcli:/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1>
ll copies
/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1/copies:
Name
-----
Pro-1
Pro-2
REP-1
REP-2
REP-3
Vplexcli:/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1>
ll links
/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1/links:
Name
-----
Pro-1->REP-1
Pro-1->REP-3
Pro-2->REP-2
Pro-2->REP-3
Vplexcli:/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1>
ll replication-sets
/recoverpoint/rpa-clusters/10.6.210.87/consistency-groups/CG-1/
replication-sets:
Name
-----
RSet0
RSet1
RSet2
RSet3

```

```

vplexcli:/> ll recoverpoint/rpa-clusters/10.6.210.87/volumes
/recoverpoint/rpa-clusters/10.6.210.87/volumes:
Name          RPA Site  RP Type  RP Role  RP  VPLEX Group  Capacity
-----
device_ty1_pro_journal_10g_1_1_vol  Tylenol-1 Journal  -        CG-3  tylenol1_rp_sys_cg  10G
device_ty1_pro_journal_10g_2_1_vol  Tylenol-1 Journal  -        CG-3  tylenol1_rp_sys_cg  10G
device_ty1_rep_5g_1_1_vol           Tylenol-1 Replication Remote
Replica      CG-3  cg1_replic          5G
device_ty1_rep_5g_2_1_vol           Tylenol-1 Replication Remote
Replica      CG-2  cg1_replic          5G
device_ty1_rep_5g_3_1_vol           Tylenol-1 Replication Remote
Replica      CG-2  cg2_rep             5G
device_ty1_rep_5g_4_1_vol           Tylenol-1 Replication Remote
Replica      CG-2  cg2_rep             5G
device_ty1_rep_5g_5_1_vol           Tylenol-1 Replication Remote
Replica      CG-1  cg2_rep             5G
device_ty1_rep_journal_10g_2_1_vol  Tylenol-1 Journal  -        CG-1  tylenol1_rp_sys_cg  10G
device_ty1_rep_journal_10g_3_1_vol  Tylenol-1 Journal  -        CG-2  tylenol1_rp_sys_cg  10G
device_ty1_rep_journal_10g_4_1_vol  Tylenol-1 Journal  -        CG-2  tylenol1_rp_sys_cg  10G
device_ty1_rep_journal_10g_5_1_vol  Tylenol-1 Journal  -        CG-1  tylenol1_rp_sys_cg  10G
device_ty1_repository_5g_1_1_vol    Tylenol-1 Repository -        -      tylenol1_rp_sys_cg  5G
ty_dr1_pro_0001_12_vol              Tylenol-1 Replication Production
Source      CG-2  cg1_pro             5G
ty_dr1_pro_0002_12_vol              Tylenol-1 Replication Production
Source      CG-3  cg2_pro             5G
ty_dr1_pro_0003_12_vol              Tylenol-1 Replication Production
Source      CG-2  cg2_pro             5G
ty_dr1_pro_0004_12_vol              Tylenol-1 Replication Production
Source      CG-2  cg2_pro             5G
ty_dr1_pro_vol                      Tylenol-1 Replication Production
Source      CG-1  cg1_pro             5G

```

```

Vplexcli:/> ll recoverpoint/rpa-clusters/10.6.210.87/volumes/
ty_dr1_pro_vol

```

```

/recoverpoint/rpa-clusters/10.6.210.87/volumes/ty_dr1_pro_vol:
Name                               Value
-----
rp-consistency-group               CG-1
rp-consistency-group-copy         Pro-1
rp-replication-set                 RSet0
rp-role                            Production Source
rp-type                            Replication
rpa-site                           Tylenol-1
size                               5G
uid                                6000144000000010f03de8cb2f4c66d9
vplex-cluster                      cluster-1
vplex-consistency-group            cg1_pro

```

See also

- `rp rpa-cluster remove`
- `rp summary`
- `rp validate-configuration`

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	
<code>[-r --rpa-cluster] <i>IP-address</i></code>	The site management IP address of the RPA cluster to remove.

Description

Removes information about a RPA cluster from VPLEX.

Removes the following commands and contexts from the VPLEX CLI:

- `/recoverpoint/rpa-clusters/ip_address/volumes`
- `/clusters/cluster_name/consistency-groups/cg_name/recoverpoint`
- `rp summary` command
- `rp validate-configuration` command

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
Vplexcli:/> ls recoverpoint/rpa-clusters/10.6.210.75
ls: No context found for 'recoverpoint/rpa-clusters/10.6.210.75'
    
```

See also

- rp rpa-cluster add
- rp summary

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
```

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.

RecoverPoint MetroPoint summary information is included in the totals.

Note: Distributed volumes used as production source volumes in MetroPoint replication will have their capacity added to both cluster 1 and cluster 2 totals. Adding the total replicated capacity of both clusters together produces a number that is larger than the actual total replicated capacity of the VPLEX Metro system, as those particular volumes are counted twice.

A summary of the number of production source volumes for MetroPoint groups is displayed after the individual cluster summary, showing the shared capacity of those volumes.

Table 21 rp summary display fields

Field	Description
VPLEX Cluster	Name of the VPLEX cluster associated with a RecoverPoint RPA.
RP Type	Type of volume. Journal - 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.

Table 21 rp summary display fields (continued)

Field	Description
	<p>Replication - All volumes of a replication set. One production volume and one or two replica volumes.</p> <p>Repository - 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. 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"> • 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. • After a permanent failover: the Replica at Production becomes the Production Source and the Remote (or Local) Source becomes the Remote (or Local) Replica. <p>Local Replica - The replica at the local site. Also, the role of the local copy, before failover.</p> <p>Production Source - Production copy, before failover.</p> <p>Remote Replica - 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 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.

Example

Display a VPLEX Metro with RecoverPoint RPAs deployed at both VPLEX clusters:

```

Vplexcli:/> rp summary
RecoverPoint Replication Totals:
-----
VPLEX Cluster  RP Type      RP Role      Total Volumes  Total Capacity
-----
cluster-1      Replication  Production Source  3              15G
                Local Replica  1              5G
                Remote Replica  3              15G
                Journal      5              50G
                Repository   1              5G
                Totals:      13             90G
VPLEX Cluster  RP Type      RP Role      Total Volumes  Total Capacity
-----
cluster-2      Replication  Production Source  2              10G
                Local Replica  2              10G
                Remote Replica  0              0G
                Journal      2              20G
                Repository   1              5G
                Totals:      7              45G

RecoverPoint MetroPoint Replication summary:
1 MetroPoint group(s) are configured with 2 Production Source volumes using a
total capacity of 10G.
Distributed Volumes used for MetroPoint replication will be counted in the
capacity of each cluster above.

```

See also

- `rp rpa-cluster add`
- `rp rpa-cluster remove`
- `rp validate-configuration`

rp validate-configuration

Validates the RecoverPoint splitter configuration.

Contexts

All contexts.

Syntax

```
rp validate-configuration
```

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 back end ports (dual-channel mode). See the second example listed below.

Best practice is to zone every RPA port to both VPLEX front end and back end 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 volume 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 storage views 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.

```

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.

```

Detect a mis-configured storage-view:

```

Vplexcli:/> rp validate-configuration
.
.

```

```

.
Validating that storage views do not have mixed non-recoverpoint and
recoverpoint initiator
ports
Storage view rp-view-demo has mixed types of initiator ports.
ERROR
.
.
=====
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
- rp rpa-cluster remove
- rp summary

schedule add

Schedules a job to run at the specified times.

Contexts

All contexts.

Syntax

schedule add

[-t|--time] *time*

[-c|--command] *command*

Arguments

Required arguments	
[-t --time] <i>time</i>	<p>* Date and time the job executes in crontab-style format enclosed in quote marks. Values are specified in the crontab-style format: “<i>Minute Hour Day-of-the-Month Month Day-of-the-week</i>”</p> <ul style="list-style-type: none"> • 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] <i>command</i>	<p>* The CLI command to be executed at the specified time.</p>

* - argument is positional.

Examples

To run the tree command every night at 1:00 a.m.:

```
VPlexcli:/> schedule add "0 1 * * *" --command tree
```

See also

- schedule list
- schedule modify
- schedule remove

schedule list

Lists all scheduled jobs.

Contexts

All contexts.

Syntax

```
schedule list
```

Examples

```
VPlexcli:/> schedule list
[0] 30 13 * * 3 syrcollect
[1] * 1 * * * tree
[2] * 2 * * * tree
```

See also

- schedule modify
- schedule remove

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

<code>[-j --job] <i>job-ID</i></code>	* ID of the scheduled job as displayed by the schedule list command.
<code>[-t --time] <i>time</i></code>	* Date and time the job executes in crontab-style format enclosed in quote marks. Values are specified in the following format: “ <i>MinuteHourDay-of-the-Month MonthDay of -he-week</i> ” 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
<code>[-c --command] <i>command</i></code>	* The CLI command to be executed at the specified time.

* - argument is positional.

Examples

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 "0 11 * * *" -c tree
```

See also

- `schedule list`
- `schedule remove`

schedule remove

Removes a scheduled job.

Contexts

All contexts.

Syntax

`schedule remove`

`[-j|--job] job-ID`

Arguments

Required arguments	
<code>[-j --job] <i>job-ID</i></code>	* 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`
- `schedule modify`

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	
<code>[-d --dayOfWeek] [0-6]</code>	Day of the week run the collection. Valid values are 0-6, where Sunday = 0...Saturday = 6.
<code>[-t --hours] [0-23]</code>	Hour at which to run the collection.
<code>[-m --minutes] [0-59]</code>	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 Dell 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 an SYR collection for every Wednesday at 12:30 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`
- `configuration event-notices-reports reset`
- `schedule list`
- `scheduleSYR list`
- `scheduleSYR remove`
- `syrcollect`

scheduleSYR list

Lists the scheduled SYR data collection job.

Contexts

All contexts.

Syntax

```
scheduleSYR list
```

Example

List the SYC collection schedule:

```
VPlexcli:/> 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`
- `configuration event-notices-reports reset`
- `scheduleSYR add`
- `scheduleSYR remove`

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`
- `configuration event-notices-reports reset`
- `scheduleSYR add`
- `scheduleSYR list`

script

Changes to interactive Jython scripting mode.

Contexts

All contexts.

Syntax

```
script
```

```
[-i|--import] module
```

```
[-u|--unimport] module
```

Arguments

Optional arguments	
--------------------	--

<code>[-i --import] <i>module</i></code>	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.
<code>[-u --unimport] <i>module</i></code>	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 or 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`

security configure-mmcs-users

Configures MMCS user accounts to synchronize service user password to the peer MMCS.

Contexts

All contexts.

Syntax

```
security configure-mmcs-users [-h | --help] [--verbose]
```

Arguments

Optional arguments	
<code>[-h --help]</code>	Displays the usage for this command.

[--verbose]	Provides more output during command execution. This may not have any effect for some commands.
-------------	--

Description

Command to sync the service account credentials between both MMCS of a V6 cluster. By running this command, password of the service account is synced to the peer MMCS.

security create-ca-certificate

Creates a new Certification Authority (CA) certificate.

Contexts

All contexts.

Syntax

```
security create-ca-certificate
  [-l|--keylength] length
  [-d|--days] days
  [-o|--ca-cert-outfilename] filename
  [-f|--ca-key-outfilename] filename
  [-s|--ca-subject-filename] filename
```

Arguments

Optional arguments	
[-l --keylength] <i>length</i>	The length (number of bits) for the CA key. Default: 2048. Range: 384 - 2048
[-d --days] <i>days</i>	Number of days that the certificate is valid. Default: 1825 (5 years). Range: 365 - 1825.
[-o --ca-cert-outfilename] <i>filename</i>	CA Certificate output filename. Default: strongswanCert.pem.
[-f --ca-key-outfilename] <i>filename</i>	CA Key output filename. Default: strongswanKey.pem.
[-s --ca-subject-filename] <i>filename</i>	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 clusters.

The system uses the Certification Authority (CA) 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

Note: The passphrase used during the VPN configuration can contain letters, numerals, and special characters.

Examples

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

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`
- `security create-host-certificate`
- `security delete-ca-certificate`
- `security delete-host-certificate`
- `security export-ca-certificate`
- `security export-host-certificate`

- security import-ca-certificate
- security import-host-certificate
- security ipsec-configure
- security show-cert-subj
- *Dell EMC VPLEX Security Configuration Guide*
- Renew security certificate procedures in the VPLEX SolVe Desktop

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] name
```

```
[-u|--org-unit] 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] <i>filename</i>	The filename of the subject file.
Optional arguments	
[-c --country] <i>country</i>	The Country value for the country key in the subject file.
[-s --state] <i>state</i>	The State value for the state key in the subject file.
[-m --org-name] <i>name</i>	Organizational Name value for the organizational name key in the subject file.
[-u --org-unit] <i>unit</i>	Organizational Unit value for the organizational unit key in the subject file.
[-l --locality] <i>locality</i>	Locality value for the locality key in the subject file.
[-n --common-name] <i>name</i>	Name value for the name key in the subject file.

<code>[-e --email] <i>e-mail</i></code>	E-mail value for the e-mail key in the subject file.
<code>--force</code>	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 <code>--force</code> argument is not specified, the command fails.

Description

Creates a subject file used in certificate creation.

Examples

Create a default certificate subject file:

```
VPlexcli:/> security create-certificate-subject --subject-out-filename
TestSubject.txt
```

The command creates a certificate subject file `TestSubject.txt` file in the `/var/log/VPlex/cli` directory with the following information:

- SUBJECT_COUNTRY=US
- SUBJECT_STATE=Massachusetts
- SUBJECT_LOCALITY=Hopkinton
- SUBJECT_ORG=Dell EMC
- SUBJECT_ORG_UNIT=Dell EMC
- SUBJECT_COMMON_NAME=FNMO0102200421
- SUBJECT_EMAIL=support@emc.com

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 command creates the certificate subject file `TestSubject.txt` file in the `/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-certificate-subject`
- `security create-host-certificate`
- `security export-ca-certificate`

- security export-host-certificate
- security import-ca-certificate
- security import-host-certificate
- security ipsec-configure
- security show-cert-subj
- *Dell 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] length

[-d|--days] days
[-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]
```

Arguments

Optional arguments	
<code>[-l --keylength] <i>length</i></code>	The length (number of bits) for the CA key. Default: 2048. Range: 384 - 2048.
<code>[-d --days] <i>days</i></code>	Number of days that the certificate is valid. Default: 730 (2 years). Range: 365 - 730.
<code>[-o --host-cert-outfilename] <i>filename</i></code>	Host certificate output filename. Default: hostCert.pem.
<code>[-f --host-key-outfilename] <i>filename</i></code>	Host key output filename. Default: hostKey.pem.
<code>[-r --host-cert-req-outfilename] <i>filename</i></code>	host certificate Request output filename. Default: hostCertReq.pem.

<code>[-c --ca-cert-filename]</code> <i>ca-certificate</i>	CA certificate used to sign the host certificate. Default: <code>strongswanCert.pem</code> .
<code>[-k --ca-key-filename]</code> <i>ca-key</i>	CA Key used to sign the host certificate. Default: <code>strongswanKey.pem</code>
<code>[-s --host-subject-filename]</code> <i>filename</i>	File that contains the subject information to create the host certificate.
<code>[-g --get-master-ca]</code>	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`.

Note: The passphrase used during the VPN configuration can contain letters, numbers, and special characters.

Examples

Create a default host certificate with the default host certificate subject information:

```
VPlexcli:/> 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`

Create a custom host certificate with the default host certificate subject information:

```
VPlexcli:/> 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`

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.

```
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-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-certificate-subject`
- `security create-host-certificate`
- `security delete-ca-certificate`
- `security delete-host-certificate`
- `security export-ca-certificate`
- `security export-host-certificate`
- `security import-ca-certificate`
- `security import-host-certificate`
- `security ipsec-configure`
- `security show-cert-subj`
- *Dell EMC VPLEX Security Configuration Guide*

security configure-certificates

Creates self-signed or imports new vendor signed certificates .

Contexts

All contexts

Syntax

`configure-certificates`

`[-p | --web-cacert-filepath=] web-cacert-filepath`

`[-k | --web-host-key-filepath=] web-host-key-filepath`

`[-w | --web-host-cert-filepath=] web-host-cert-filepath`

`[-n | --vpn-cacert-filepath=] vpn-cacert-filepath`

`[-m | --vpn-host-key-filepath=] vpn-host-key-filepath`

`[-c | --vpn-host-cert-filepath=] vpn-host-cert-filepath`

`--verbose`

Arguments

<code>[-p --web-cacert-filepath=] <i>web-cacert-filepath</i></code>	Specifies the Web CA certificate file path.
<code>[-k --web-host-key-filepath=] <i>web-host-key-filepath</i></code>	Specifies the path to Web host key file.
<code>[-w --web-host-cert-filepath=] <i>web-host-cert-filepath</i></code>	Specifies the path to Web host certificate file.
<code>[-n --vpn-cacert-filepath=] <i>vpn-cacert-filepath</i></code>	Specifies the host CA certificate file path.
<code>[-m --vpn-host-key-filepath=] <i>vpn-host-key-filepath</i></code>	Specifies the VPN host key file path.
<code>[-c --vpn-host-cert-filepath=] <i>vpn-host-cert-filepath</i></code>	Specifies the path to the VPN host certificate file.
<code>--verbose</code>	Provides more output during command execution. This may not have any effect for some commands.
<code>[-h --help]</code>	Displays the usage for this command.

Description

A management server authenticates an external client entity based on the Certificate Authority it trusts. The trust store is used for Web/ REST clients over HTTPS connections and an SSL database for inter-site VPN connections. The CA trust can be self-signed based on a local certificate authority subject info or a third-party vendor (such as Verisign, Globalsign) signed. A plain run of the command without any options creates self-signed CA certificates. This also creates host and web certificates which are signed by the self signed CA certificate created. Running the command with options `--vpn-host-cert-filepath`, `--vpn-host-key-filepath` and `--vpn-cacert-filepath` imports the vendor signed certificates provided to configure a VPN between the sites and the VPLEX Witness server. This also configures web with self-signed CA/web certificates. Running the command with options `--web-host-cert-filepath`, `--web-host-key-filepath` and `--web-cacert-filepath` imports the vendor signed certificates provided to configure web. This also configures vpn with self-signed CA/ host certificates. If all options `--vpn-host-cert-filepath`, `--vpn-host-key-filepath`, `--vpn-cacert-filepath`, `--web-host-cert-filepath`, `--web-host-key-filepath`, and `--web-cacert-filepath` are used then the command configures both vpn and web with imported vendor signed certificates provided. For self-signed certificates the default values of validity days and keylength are used. This command is the same as running `security create-ca-certificate` and `security create-host-certificate` 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 passphrase used during the VPN configuration can contain letters, numbers, and special characters.

Examples

To create self-signed certificates for VPLEX Local configurations

1. Use the `security configure-certificates` command to create self-signed CA/host certificates for both vpn and web.
2. Enter and write down the passphrase entered during the creation of CA certificate for vpn configuration.

3. Enter and write down the passphrase entered during the creation of host certificate for vpn configuration.
4. Enter and write down the passphrase entered during the creation of CA certificate for web configuration.
5. Enter and write down the passphrase entered during the creation of host certificate for web configuration.

```

Vplexcli:/> security configure-certificates
Default self-signed host certificates will be generated. Default self-signed
web certificates will be generated.
Please create a passphrase (at least 8 characters) for the Certificate
Authority Key. Make a note of this passphrase as you
will need it to configure a second cluster of a VPLEX Metro or Geo.
Certificate Authority passphrase:
Re-enter the passphrase for the Certificate Authority Key:
New CA certificate /etc/ipsec.d/cacerts/strongswanCert.pem created New CA
key /etc/ipsec.d/private/strongswanKey.pem
created
Please create a passphrase (at least 8 characters) for the Local Host
Certificate Keys to be used to configure the VPN. Make a note
of this passphrase as you will need it later. Host Certificate passphrase
Re-enter:
New Host certificate request /etc/ipsec.d/reqs/hostCertReq.pem created New
Host certificate /etc/ipsec.d/certs/hostCert.pem created and signed by the CA
Certificate /etc/ipsec.d/cacerts/strongswanCert.pem
Please create a passphrase (at least 8 characters) for the Web Host
Certificate Keys to be used to configure the web server. Make a note of this
passphrase as you will need it later. Web Host Certificate passphrase:
Re-enter:
New Host certificate request /etc/ipsec.d/reqs/webServerHostCertReq.pem
created New Host certificate /etc/ipsec.d/certs/webServerHostCertFile.pem
created and signed by the CA Certificate /etc/ipsec.d/cacerts/strongswanCert.pem
    
```

To import vendor-signed certificates on a VPLEX Local:

1. Use the `security configure-certificates` command with options `--web-host-cert-filepath`, `--web-host-key-filepath` and `--web-cacert-filepath` to import the vendor signed certificates for web configuration.
2. Enter the passphrase (if present) for the imported CA certificate for web configuration.
3. Enter the passphrase (if present) for the imported host certificate for web configuration.
4. Enter and write down the passphrase entered during the creation of self-signed CA certificate for vpn configuration.
5. Enter and write down the passphrase entered during the creation of self-signed host certificate for vpn configuration.

```

Vplexcli:/> security configure-certificates -w /tmp/certs/hostCert.pem -
k /tmp/certs/hostKey.pem -p /tmp/certs/strongswanCert.pem

Default self-signed host certificates will be generated.

Input file is a root CA but the CA signer not found in local ssl database
Added imported CA signer to local certificate database
Please enter the passphrase for the imported web certificate key (at least
8characters) :
Re-enter:
CA Certificate /tmp/certs/strongswanCert.pem successfully imported

Host Certificate /tmp/certs/hostCert.pem successfully imported Host Key /tmp/
certs/hostKey.pem successfully imported
    
```

```
Please create a passphrase (at least 8 characters) for the Certificate Authority Key. Make a note of this passphrase as you will need it to configure a second cluster of a VPLEX Metro or Geo. Certificate Authority passphrase:
```

```
Re-enter the passphrase for the Certificate Authority Key:
```

```
New CA certificate /etc/ipsec.d/cacerts/strongswanCert.pem created New CA key /etc/ipsec.d/private/strongswanKey.pem created
```

```
Please create a passphrase (at least 8 characters) for the Local Host Certificate Keys to be used to configure the VPN. Make a note of this passphrase as you will need it later.
```

```
Host Certificate passphrase:
```

```
Re-enter:
```

```
New Host certificate request /etc/ipsec.d/reqs/hostCertReq.pem created New Host certificate /etc/ipsec.d/certs/hostCert.pem created and signed by the CA Certificate /etc/ipsec.d/cacerts/strongswanCert.pem
```

To configure certificates on a VPLEX Metro:

i Note: To import vendor-signed certificates for a VPLEX Metro: the VPN host certificates to be imported in both VPLEX clusters should be signed and issued by the same CA. Similarly, the Web host certificates to be imported in both VPLEX clusters should be signed and issued by the same CA.

i 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 passphrase used during the VPN configuration can contain letters, numbers, and special characters.

1. Follow the procedure for Local configuration on cluster 1.

2. On cluster 2, Do one of the following:

- For self-signed certificates
Similar to configuring self-signed certificates on first cluster, but specify the IP address and service account password of the first cluster.
- For vendor-signed certificates
Similar to configuring vendor signed certificates on first cluster, but specify the IP address and service account password of the first cluster.

```
Vplexcli: /> security configure-certificates -w /tmp/certs/hostCert.pem -k /tmp/certs/hostKey.pem -p /tmp/certs/strongswanCert.pem
Default self-signed host certificates will be generated.
Please enter the IP address of the VPLEX cluster on which the original CA was created: 10.110.19.37
Please enter the 'service' account password( 8 chars ) :
Re-enter the password:
Input file is a root CA but the CA signer not found in local ssl database Added imported CA signer to local certificate database
Please enter the passphrase for the imported web certificate key (at least 8 characters) :
Re-enter: CA Certificate /tmp/certs/strongswanCert.pem successfully imported
Host Certificate /tmp/certs/hostCert.pem successfully imported Host Key /tmp/certs/hostKey.pem successfully imported
Please create a passphrase (at least 8 characters) for the Certificate Authority Key. Make a note of this passphrase as you will need it to configure a second cluster of a VPLEX Metro or Geo. Certificate Authority passphrase:
Re-enter the passphrase for the Certificate Authority Key:
New CA certificate /etc/ipsec.d/cacerts/strongswanCert.pem created New CA key /etc/ipsec.d/private/strongswanKey.pem created
Please create a passphrase (at least 8 characters) for the Local Host
```

```
Certificate Keys to be used to configure the VPN. Make a note of this
passphrase as you will need it later.
Host Certificate passphrase:
Re-enter:
New Host certificate request /etc/ipsec.d/reqs/hostCertReq.pem created
New Host certificate /etc/ipsec.d/certs/hostCert.pem created and signed
by the CA Certificate /etc/ipsec.d/cacerts/strongswanCert.pem
```

See also

- security list-certificates
- security web-configure
- configuration cw-vpn-configure

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

Optional arguments	
<code>[-o --ca-cert-outfilename] <i>filename</i></code>	CA Certificate output filename. Default: strongswanCert.pem.
<code>[-f --ca-key-outfilename] <i>filename</i></code>	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.

Examples

Delete a custom CA certificate (not the default):

```
VPlexcli:/> security delete-ca-certificate -o TestCACert.pem -f TestCAKey.pem
```

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

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`
- `security delete-host-certificate`
- *Dell EMC VPLEX Security Configuration Guide*

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

Optional arguments	
<code>[-o --host-cert-outfilename] <i>filename</i></code>	host certificate output filename. Default: hostCert.pem.
<code>[-f --host-key-outfilename] <i>filename</i></code>	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.

Examples

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

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
- security create-host-certificate
- security delete-ca-certificate

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] filepath
[-k|--ca-key-filepath] filepath
[-e|--ca-export-location] path
```

Arguments

Required arguments	
[-e --ca-export-location] <i>path</i>	The absolute path of the location to which to export the CA Certificate and CA Key.
Optional arguments	
[-c --ca-cert-filepath] <i>filepath</i>	The absolute path of the CA certificate file to export. Default: /etc/ipsec.d/cacerts/strongswanCert.pem.
[-k --ca-key-filepath] <i>filepath</i>	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:** You must have write privileges at the location to which you export the certificate.

The import or 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
```


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`
- `security create-certificate-subject`
- `security create-host-certificate`
- `security export-host-certificate`
- `security import-ca-certificate`
- `security import-host-certificate`
- `security ipsec-configure`
- `security show-cert-subj`
- *Dell 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] path
[-k|--host-key-filepath] path
[-e|--host-export-location] path
```

Arguments

Required arguments	
<code>[-e --host-export-location] <i>path</i></code>	The absolute path of the location to which to export the host certificate and host key.
Optional arguments	
<code>[-c --host-cert-filepath] <i>path</i></code>	The absolute path of the host certificate file to export. Default: <code>/etc/ipsec.d/certs/hostCert.pem</code>
<code>[-k --host-key-filepath] <i>path</i></code>	The absolute path of the host key file to export. Default: <code>/etc/ipsec.d/private/hostKey.pem</code>

Description

Exports the host certificate to the specified location.

 **Note:** You must have write privileges at the location to which you export the certificate .

Example

Export the default host certificate and key to `/var/log/VPlex/cli`:

```
VPlexcli:/> security export-host-certificate -e /var/log/VPlex/cli
```

Export a custom host certificate and its 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`
- `security create-certificate-subject`
- `security create-host-certificate`
- `security export-ca-certificate`
- `security import-ca-certificate`
- `security show-cert-subj`
- *Dell EMC VPlex Security Configuration Guide*

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] path
[-k|--ca-key-filepath] path
[-i|--ca-cert-import-location] location
[-j|--ca-key-import-location] location
```

Arguments

Required arguments	
<code>[-c --ca-cert-filepath] <i>path</i></code>	The absolute path of the CA certificate file to import.
<code>[-k --ca-key-filepath] <i>path</i></code>	The absolute path of the CA key file to import.
Optional arguments	
<code>[-i --ca-cert-import-location] <i>location</i></code>	The absolute path of the location to which to import the CA certificate. Default location - <code>/etc/ipsec.d/cacerts</code> .
<code>[-j --ca-key-import-location] <i>location</i></code>	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: You must have write privileges at the location from which you import the certificate.

If the import locations for the CA certificate and CA key has files with the same names, they are overwritten.

The import or 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

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
- security create-certificate-subject
- security create-host-certificate
- security export-ca-certificate
- security import-host-certificate
- security ipsec-configure
- security show-cert-subj
- *Dell EMC VPLEX Security Configuration Guide*

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] path`

`[-k|--host-key-filepath] path`

`[-i|--host-cert-import-location] path`

`[-j|--host-key-import-location] path`

Arguments

Required arguments	
<code>[-c --host-cert-filepath] <i>path</i></code>	The absolute path of the host certificate file to import.
<code>[-k --host-key-filepath] <i>path</i></code>	The absolute path of the host key file to import.
Optional arguments	
<code>[-i --host-cert-import-location] <i>path</i></code>	The absolute path of the location to which to import the host certificate. Default location - <code>/etc/ipsec.d/certs</code> .
<code>[-j --host-key-import-location] <i>path</i></code>	The absolute path of the location to which to import the host certificate. Default location - <code>/etc/ipsec.d/private</code> .

Description

Imports the host certificate from the specified location.

i **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.

Examples

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`

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`
- `security create-certificate-subject`
- `security create-host-certificate`
- `security export-ca-certificate`
- `security import-ca-certificate`
- `security ipsec-configure`
- `security show-cert-subj`
- *Dell EMC VPLEX Security Configuration Guide*

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-address
[ -c | --host-cert-filename ] host-certificate
[ -k | --host-key-filename ] host-key
```

Arguments

Required arguments	
[-i --remote-ms-ipaddr] <i>remote-address</i>	IP address of the remote management server.
Optional arguments	
[-c --host-cert-filename] <i>host-certificate</i>	host certificate filename.
[-k --host-key-filename] <i>host-key</i>	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 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 IP seed is used to generate the IP addresses used by the internal components of the system. For more information about components and their IP addresses, refer to the *EMC VPLEX Configuration 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).

Examples

In the following example:

- On first cluster: the `security ipsec-configure` command configures IPsec to second cluster at 10.6.209.33

```
Vplexcli:/> security ipsec-configure -i 10.6.209.33
```

- On second cluster: the `security ipsec-configure` command configures IPsec to first cluster at 10.6.209.32

```
Vplexcli:/> security ipsec-configure -i 10.6.209.32
```

- On the first cluster, the `vpn status` command confirms that the VPN to the second cluster is up

```
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
```

- On the second cluster, the `vpn status` command confirms that the VPN to the first cluster is up

```
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`
- `security create-certificate-subject`
- `security create-host-certificate`
- `security export-ca-certificate`
- `security import-ca-certificate`
- `security show-cert-subj`
- *Dell EMC VPLEX Security Configuration Guide*

security list-certificates

Displays the validation status of the existing Certificates

Contexts

All contexts

Syntax

```
security list-certificates
```

```
[-h | --help]
```

```
[--verbose]
```

Arguments

<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

The command lists down all the certificates present in the system along with the validation status of each parameter associated with the certificate.

The following parameters are currently displayed :

Table 22 Certificate parameters

Filename	Name of the certificate file
Cert Type	Type of the certificate file. This column tells if a certificate is a host or a root certificate Possible values are : CA for a root certificate Host for a host certificate

Table 22 Certificate parameters (continued)

Issued to	Displays the Common Name of the Subject
Issued by	Displays the Common Name of the Issuer
Date Validity	Displays the date validation status of a certificate Possible values are : YES, when a certificate is valid EXPIRED, when a certificate is expired i.e current date is post expiry date NOT YET VALID, when a certificate is not yet valid i.e certificate start date is post current date
Expiry Date	Displays the expiry date of a certificate
Start Date	Displays the start date of a certificate
Issuer check	Displays the validation status of signer of the certificate Possible values are : YES, when the certificate is signed by the CA present in the system NO, when there is a mismatch in the signer of the certificate and the CA present in the system
Entry in TS/KS	Displays where the certificate entry is found. Possible values are : TS, when the certificate has its entry in Trust Store. For example, /var/log/VPlex/cli/default.cacerts KS, when the certificate has its entry in Key Store /var/log/VPlex/cli/.keystore
Signature Algorithm	Displays the Signature Algorithm of the certificate.

See Also

- security configure-certificates
- security web-configure
- configuration cw-vpn-configure

security remove-login-banner

Removes the login banner from the management server.

Contexts

All contexts.

Syntax

```
security remove-login-banner
[-f|--force]
```


Arguments

Optional arguments	
<code>[-f --force]</code>	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`
- *Dell EMC VPLEX Security Configuration Guide*

security renew-all-certificates

Renews CA and host security certificates.

Contexts

All contexts.

Syntax

```
security renew-all-certificates
```

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

Use the `security renew-all-certificates` command to renew all security certificates on a VPLEX system.

In Metro 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 the “Before you begin” section below 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. 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.
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.

You are always prompted for the Certificate Authority (CA) passphrase when you run the command on the second cluster.

When renewing the certificates on the second cluster, you might be prompted to enter the service password. Contact the System Administrator to obtain the current service password.

- Renew the certificates using a common passphrase.
All certificates are renewed using the same passphrase.

CAUTION In Metro 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.

CAUTION

Before you begin:

1. 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.
2. 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.
3. 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.

Examples

The following is the simplest renewal case. It is the case of renewing a certificate on a 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

```

An example of running the certificate renewal on a 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

```

This example shows running the `renew-all-certificates` command 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: emc12345
Re-enter the password: emc12345
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...

```

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

- `security create-ca-certificate`
- `security create-host-certificate`
- `security export-ca-certificate`
- `security import-ca-certificate`
- `security import-host-certificate`

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	
<code>[-b --login-banner-file] <i>file</i></code>	Full pathname to the file containing the formatted login banner text.
Optional arguments	
<code>[-f --force]</code>	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.

Examples

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
- EMC VPLEX Security Configuration Guide

security show-cert-subj

Displays the certificate subject file.

Contexts

All contexts.

Syntax

```
security show-cert-subj
[s|--subject-infilename] filename
```

Arguments

Required arguments	
<code>[-s --subject-infilename] <i>filename</i></code>	Filename of the certificate subject file to display. The file is assumed to reside in the following directory on the management server: <code>/var/log/VPlex/cli</code>

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=FNM00094400134
SUBJECT_EMAIL=support@emc.com
```

See also

- security create-certificate-subject

security web-configure

Configures the webserver CA, certificate and key by deleting the previous entries from keystore and truststore and registering the new entries in the truststore and keystore.

Contexts

All contexts

Syntax

```
security web-configure
```

```
[-h | --help]
```

```
[--verbose]
```

Arguments

<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

Description

Use this command to configure the web certificates after importing the CA and host certificate and keys.

- Use the `security configure-certificates` command to create or import the CA and host certificate and keys to configure web.
- Run the `security web-configure` command.

You can supply the file names for all the three certificates as parameters. This command supports external certificates.

Examples

```
VPlexcli:/> security web-configure

Your Java Key Store has been created. addWebapp(/smsflex,/opt/emc/VPlex/
tomcat/webapps/smsflex) addWebapp(/vplex,/opt/emc/VPlex/tomcat/webapps/vplex)
addWebapp(/cimom,/opt/emc/VPlex/tomcat/webapps/cimom) addWebapp(/,/opt/emc/
VPlex/tomcat/webapps/ROOT) https keystore: /var/log/VPlex/cli/.keystore
started web server on ports {'http': 49880, 'https': 49881}
```

See also

- `security configure-certificates`
- `security list-certificates`
- `configure cw-vpn-configure`

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 writable attributes in the given context.

Contexts

All contexts.

Syntax

```
set
[-d|--default]
[-f|--force]
[-a|--attributes] pattern
[-v|--value] value
```

Arguments

Optional arguments	
<code>[-d --default]</code>	Sets the specified attributes to the default values, if any exist. If no attributes are specified, displays the default values for attributes in the current/specified given context.
<code>[-f --force]</code>	Force the value to be set, bypassing any confirmations or guards.
<code>[-a --attributes] <i>pattern</i></code>	* Attribute selector pattern.
<code>[-v --value] <i>value</i></code>	* The new value to assign to the specified attributes.

* - 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.

Examples

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.
.
..
```

Change the name of a meta-volume:

```
VPlexcli:/clusters/cluster-1/system-volumes/
new_meta1_backup_2010May24_163810> set name backup_May24_pre_refresh
```

Display information about attributes in the eth0 context:

```
VPlexcli:/> set /management-server/ports/eth0
attribute          input-description
-----
/management-server/ports/eth0::address  Read-only.
/management-server/ports/eth0::auto-negotiate  Read-only.
/management-server/ports/eth0::duplex  Read-only.
/management-server/ports/eth0::gateway  Read-only.
/management-server/ports/eth0::inet6-address  Read-only.
/management-server/ports/eth0::inet6-gateway  Read-only.
/management-server/ports/eth0::name  Read-only.
```

```
/management-server/ports/eth0::net-mask      Read-only.
/management-server/ports/eth0::speed        Read-only.
/management-server/ports/eth0::status       Read-only.
```

Display the address attribute for eth0:

```
VPlexcli:/> set /management-server/ports/eth0::address
attribute                                     input-description
-----
/management-server/ports/eth0::address      Read-only.
```

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
```

Attach a rule-set to cluster1_Active to the device dd_00:

```
VPlexcli:/distributed-storage/distributed-devices> set dd_00::rule-set-name
cluster1_Active
```

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                               VPD83T3:6006016061212e00b0171b696696e211
thin-rebuild                            true
total-free-space                        0B
underlying-storage-block-size           512
use                                      used
used-by                                  [extent_test01_1]
vendor-specific-name                    DGC
vias-based                               false
```

Use the `set enabled false --force` command in the `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
```

Enable and display ports:

```

Vplexcli:/engines/engine-1-1/directors/director-1-1-A/hardware/ports> ll
Name          Address          Role             Port Status
-----
A0-FC00       0x5000144260006e00 front-end        no-link
A0-FC01       0x5000144260006e01 front-end        up
A0-FC02       0x5000144260006e02 front-end        up
A0-FC03       0x0000000000000000 front-end        down
A1-FC00       0x5000144260006e10 back-end         up
A1-FC01       0x5000144260006e11 back-end         up
A1-FC02       0x5000144260006e12 back-end         no-link
A1-FC03       0x5000144260006e13 back-end         no-link
A2-FC00       0x5000144260006e20 wan-com          up
A2-FC01       0x5000144260006e21 wan-com          up
A2-FC02       0x5000144260006e22 wan-com          no-link
A2-FC03       0x5000144260006e23 wan-com          no-link
A3-FC00       0x5000144260006e30 local-com        up
A3-FC01       0x5000144260006e31 local-com        up
A3-FC02       0x0000000000000000 -                down
A3-FC03       0x0000000000000000 -                down
Vplexcli:/engines/engine-1-1/directors/director-1-1-A/hardware/ports> set A0-FC03::enabled true
Vplexcli:/engines/engine-1-1/directors/director-1-1-A/hardware/ports> ll
Name          Address          Role             Port Status
-----
A0-FC00       0x5000144260006e00 front-end        no-link
A0-FC01       0x5000144260006e01 front-end        up
A0-FC02       0x5000144260006e02 front-end        up
A0-FC03       0x5000144260006e03 front-end        no-link
A1-FC00       0x5000144260006e10 back-end         up
A1-FC01       0x5000144260006e11 back-end         up
A1-FC02       0x5000144260006e12 back-end         no-link
A1-FC03       0x5000144260006e13 back-end         no-link
A2-FC00       0x5000144260006e20 wan-com          up
A2-FC01       0x5000144260006e21 wan-com          up
A2-FC02       0x5000144260006e22 wan-com          no-link
A2-FC03       0x5000144260006e23 wan-com          no-link
A3-FC00       0x5000144260006e30 local-com        up
A3-FC01       0x5000144260006e31 local-com        up
A3-FC02       0x0000000000000000 -                down
A3-FC03       0x0000000000000000 -                down

```

Change and display the name of a virtual volume:

i **Note:** Changing a virtual volume name will not cause any impact to host I/O.

```

Vplexcli:/clusters/cluster-1/virtual-volumes/EMC-CLARiiON-0075-VNX-
LUN122_1_vol>set -a name -v new_name
Vplexcli:/clusters/cluster-1/virtual-volumes/new_name> ll
Name          Value
-----
block-count   2621440
block-size    4K
cache-mode    synchronous
capacity      10G
consistency-group -
expandable    true
health-indications []
health-state  ok
locality      local
operational-status ok
scsi-release-delay 0
service-status running

```

```
storage-tier -
supporting-device device EMC-CLARiiON-APM00113700075-VNX_LUN122_1
system-id EMC-CLARiiON-0075-VNX-LUN122_1_vol
volume-type virtual-volume
```

Return to the virtual-volumes context and change directory to the new name:

```
VPlexcli:/clusters/cluster-1/virtual-volumes/new_name> cd ..
VPlexcli:/clusters/cluster-1/virtual-volumes> cd new_name
```

Run a listing on the volume to display the new name for the system-id:

```
VPlexcli:/clusters/cluster-1/virtual-volumes/new_name> ll
Name Value
-----
block-count 2621440
block-size 4K
cache-mode synchronous
capacity 10G
consistency-group -
expandable true
health-indications []
health-state ok
locality local
operational-status ok
scsi-release-delay 0
service-status running
storage-tier -
supporting-device device EMC-CLARiiON-APM00113700075-VNX_LUN122_1
system-id new_name
volume-type virtual-volume
```

Set the SPC version to Version 3 on an initiator port:

```
VPlexcli:/clusters/cluster-1/exports/initiator-ports/test_port_1> set scsi-
scs-version 3
To avoid Data Unavailability and host issues, follow the procedure described
in the KB article 'SPC-3 support in VPLEX' to change the SPC version.
SPC-3 must be applied only on the supported Operating Systems that are listed
in the KB article. Continue? (Yes/No)
```

```
VPlexcli:/clusters/cluster-1/exports/initiator-ports/test_port_1> ll
Name Value
-----
node-wwn 0x20000025b505003f
port-wwn 0x200000cc05bb002e
scsi-scs-version 3
suspend-on-detach -
target-ports [P0000000043E00BDD-A0-FC00, P0000000043E00BDD-A0-FC01,
P0000000043F00BDD-B0-FC00, P0000000043F00BDD-B0-FC01]
type default
```

Set the SPC version to Version 3 on a storage view:

```
VPlexcli:/clusters/cluster-1/exports/storage-views/test_view_1> set scsi-scs-
version 3
To avoid Data Unavailability and host issues, follow the procedure described
```

```
in the KB article 'SPC-3 support in VPLEX' to change the SPC version.
SPC-3 must be applied only on the supported Operating Systems that are listed
in the KB article. The new SPC version is applied to all the initiators
in the storage-view. Continue? (Yes/No) Yes
```

```
VPlexcli:/clusters/cluster-1/exports/storage-views/test_view_1> ll
Name                               Value
-----
-----
caw-enabled                         true
controller-tag                      -
initiators                         [test_port]
operational-status                  ok
port-name-enabled-status            [P0000000043E00BDD-A0-FC00,true,ok,
P0000000043E00BDD-A0-FC01,true,ok,
P0000000043F00BDD-B0-FC00,true,ok,
P0000000043F00BDD-B0-FC01,true,ok]
ports                               [P0000000043E00BDD-A0-FC00, P0000000043E00BDD-A0-
FC01, P0000000043F00BDD-B0-FC00,
P0000000043F00BDD-B0-FC01]
scsi-spc-version                    3
virtual-volumes                     [(0,device_C1-
RHEL_XtremIO0547_LUN_00001_1_vol,VPD83T3:600014400000010f00bddd268733d19,200G
)]
write-same-16-enabled               true
xcopy-enabled                       true
```

See also

- storage-volume claim
- storage-volume unclaim

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.

Note: According to best practices, the front-end ports should be set to the default p2p and connected to the hosts via a switched fabric.

WARNING It is not recommended to change the topology on the local COM ports, as it can lead to the directors going down and data unavailability.

Example

Navigate to a Fibre Channel 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

show-use-hierarchy

Display the complete usage hierarchy for a storage element from the top-level element down to the storage-array.

Contexts

All contexts.

Syntax

```

show-use-hierarchy
[-t|--targets] path, path,...
    
```

Arguments

Required arguments	
[-t --targets] path, path,...	<p>* Comma separated list of target storage elements. You can specify meta, logging and virtual volumes, local and distributed devices, extents, storage-volumes or logical-units on a single command line.</p> <p>Note: A complete context path to the targets must be specified. For example:</p>

```
show-use-hierarchy /clusters/cluster-1/storage-
elements/storage-volumes/volume
```

or:

```
show-use-hierarchy /clusters/cluster-1/storage-
elements/storage-volumes/*
```

* - argument is positional.

Description

This command drills from the specified target up to the top-level volume and down to the storage-array. The command will detect sliced elements, drill up through all slices and indicate in the output that slices were detected. The original target is highlighted in the output.

See also

- `drill-down`
- `tree`

sms dump

Collects the logs files on the management server.

Contexts

All contexts.

Syntax

```
sms dump
```

```
[-d|--destination-directory] directory
```


```
[-t|--target_log] logName
```

Arguments

Required arguments	
<code>[-d --destination-directory] <i>directory</i></code>	Destination directory for the sms dump logs.
Optional arguments	
<code>[-t --target_log] <i>logName</i></code>	Collect only files specified under logName from smsDump.xml.

Description

Collects the following log files:

 **Note:** The log files listed below are the core set of files along with other files that are not listed.

Clilogs

- `/var/log/VPlex/cli/client.log*` -- VPlexcli logs, logs dumped by VPlexcli scripts
- `/var/log/VPlex/cli/session.log*` -- what the user does in a VPlexcli 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 VPLEXcli
- /var/log/VPLEX/cli/connections -- what the VPLEXcli 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.omsf
- /var/log/firewall
- /etc/sysconfig/SuSEfirewall2
- /etc/sysconfig/network/ifcfg*
- /etc/sysconfig/network/ifroute*
- /etc/sysctl.conf

Examples

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
- collect-diagnostics
- director appdump
- getsysinfo

snmp-agent configure

Configures the VPLEX SNMP agent service on the local cluster.

Contexts

All contexts.

Syntax

```
snmp-agent configure
```

Arguments

Optional arguments	
<code>[-s --snmp-protocol-version] <i>snmpv2c</i></code>	Configures the VPLEX SNMP agent service version 2c
<code>[-s --snmp-protocol-version] <i>snmpv3</i></code>	Configures the VPLEX SNMP agent service version 3

Description

Configures the SNMP agent on the local cluster, and starts the SNMP agent. You can configure VPLEX SNMP agent service version 2c or version 3. If you do not specify a version, the SNMP agent service version 3 is configured.

`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.

- i** **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 versions SNMPv3 and 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.

See also

- `snmp-agent start`
- `snmp-agent status`
- `snmp-agent stop`
- `snmp-agent unconfigure`

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`
- `snmp-agent status`
- `snmp-agent stop`
- `snmp-agent unconfigure`

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
- snmp-agent start
- snmp-agent stop
- snmp-agent unconfigure

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`
- `snmp-agent start`
- `snmp-agent status`
- `snmp-agent unconfigure`

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`
- `snmp-agent start`
- `snmp-agent status`
- `snmp-agent stop`

source

Reads and executes commands from a script.

Contexts

All contexts.

Syntax

`source`

`[-f|--file] filename`

Arguments

Required arguments	
<code>[-f --file] filename</code>	* 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.

Examples

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`

storage-tool dismantle

Dismantles virtual-volumes, devices (local or distributed) and extents down to the storage-volumes, including unclaiming the storage-volumes.

Contexts


All contexts.

Syntax

```
storage-tool dismantle
[--do-not-unclaim]
[-h | --help]
[--verbose]
[-f | --force]
[-s | --storage-extents= storage-extent [, storage-extent] ...]
```

Arguments

Optional arguments	
--------------------	--


<code>[-h --help]</code>	Displays the usage for this command.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.
<code>[-f --force]</code>	Do not ask for confirmation.
<code>--do-not-unclaim</code>	Skips unclaiming of the storage-volume.  Note: By default, the storage volumes are unclaimed.
Required arguments	
<code>[-s --storage-extents= <i>storage-extent</i> [, <i>storage-extent</i>] ...]</code>	Specifies the storage-extents (virtual-volumes, local or distributed devices or extents) to dismantle.

* argument is positional

Description

Dismantles virtual-volumes, devices (local or distributed) and extents down to the storage-volumes, including unclaiming the storage-volumes.

Run `storage-tool dismantle` against top-level storage elements only. If you run `storage-tool dismantle` against virtual-volumes, they must not belong to either a consistency-group or storage-view.

 **Note:** This command does NOT allow dismantling of consistency groups or storage views, or of storage extents that are not root nodes in a storage hierarchy (i.e. targets must not be supporting other storage).

The command fails with an exception before dismantling anything if:

- A volume to be dismantled is exported in a view and that view is not a dismantle target.
- A volume to be dismantled is in a consistency group and that consistency group is not a dismantle target.
- The dismantle target is supporting other storage (i.e. has anything above it).

storage-tool compose

Creates a virtual-volume on top of the specified storage-volumes, building all intermediate extents, local, and distributed devices as necessary.

Contexts


All contexts.

Syntax

```
storage-tool compose
[-n|--name] name
[-g|--geometry] {raid-0|raid-1|raid-c}
[-d|--storage-volumes] storage-volume [, storage-volume...]
[-m|--source-mirror] source-mirror
[-c|--consistency-group] consistency-group
[-v|--storage-views] storage-view [, storage-view ...]
[-t|--thin]
[-h|--help]
```

[--verbose]

Arguments

Required arguments	
<code>[-n --name] <i>name</i></code>	* Specifies the name for the new virtual volume. Must be unique across the system.
<code>[-g --geometry] {raid-0 raid-1 raid-c}</code>	* Specifies the geometry to use for the local devices at each cluster. Valid values are raid-0, raid-1, or raid-c.
Optional arguments	
<code>[-d --storage-volumes] <i>storage-volume</i> [, <i>storage-volume...</i>]</code>	* Specifies a list of storage volumes to build the virtual volume from. These may be claimed, but must be unused.
<code>[-m --source-mirror] <i>source-mirror</i></code>	Specifies the storage volume to use as a source mirror when creating local and distributed devices.  Note: If specified, <code>--source-mirror</code> will be used as a source-mirror when creating local and distributed RAID 1 devices. This will trigger a rebuild from the source-mirror to all other mirrors of the RAID 1 device (local and distributed). While the rebuild is in progress the new virtual volume (and supporting local and/or distributed devices) will be in a degraded state, which is normal. This option only applies to RAID 1 local or distributed devices. The <code>--source-mirror</code> may also appear in <code>--storage-volumes</code> .
<code>[-c --consistency-group] <i>consistency-group</i></code>	Specifies the context path of a consistency group that the new virtual volume should be added to. The new virtual-volume's global geometry must be compatible with the consistency group's <code>storage-at-clusters</code> attribute.
<code>[-v --storage-views] <i>storage-view</i> [, <i>storage-view...</i>]</code>	Specifies the context path of the storage views that the new virtual volume will be added to. The new virtual volume's global geometry must be compatible with the storage view's locality.
<code>[-t --thin]</code>	Specifies whether the new virtual-volume is thin-enabled or not. The supporting storage-volumes must be thin-capable in order for a virtual-volume to be thin-enabled. The virtual-volume must also have a valid RAID geometry to be thin-enabled.
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more help during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

This command supports building local or distributed (i.e., distributed RAID 1 based) virtual volumes with RAID 0, RAID 1, or RAID C local devices. It does not support creating multi-device storage hierarchies (such as a RAID 1 on RAID 0s on RAID Cs).

For RAID 1 local devices, a maximum of eight legs may be specified.

If the new virtual volume's global geometry is not compatible with the specified consistency group or storage views, the virtual volume will not be created. However, failure to add the new virtual

volume to the specified consistency group or storage views does not constitute an overall failure to create the storage and will not be reported as such.

Note: In the event of an error, the command will not attempt to perform a roll-back and destroy any intermediate storage objects it has created. If cleanup is necessary, use the `show-use-hierarchy` command on each storage volume to identify all residual objects and delete each one manually.

The `--stop-at` option imposes the following constraints on other options:

- If `--stop-at=virtual-volume`, only the `--consistency-group` and `--storage-views` options can be specified.
- If `--stop-at=local-device`, storage-volumes from only one cluster can be specified.
- If `--stop-at=distributed-device`, storage-volumes from at least two clusters must be specified.

Example

Create a thin-capable virtual volume with RAID 0 local devices and specified storage volumes:

```

VPlexcli:/clusters/cluster-1/virtual-volumes> storage-tool compose --name
myVolume --storage-volumes thin_capable_sv_1
--thin --geometry raid-0
Successfully created /clusters/cluster-1/virtual-volumes/myVolume.
VPlexcli:/clusters/cluster-1/virtual-volumes> ll myVolume/

/clusters/cluster-1/virtual-volumes/myVolume:
Name                               Value
-----
block-count                         2621440
block-size                           4K
cache-mode                           synchronous
capacity                             10G
consistency-group                    -
expandable                           true
expandable-capacity                  0B
expansion-method                     storage-volume
expansion-status                     -
health-indications                   []
health-state                         ok
locality                             local
operational-status                   ok
recoverpoint-protection-at           []
recoverpoint-usage                   -
scsi-release-delay                   0
service-status                       unexported
storage-tier                         -
supporting-device                    device_myVolume_c1
system-id                            myVolume
thin-capable                         true
thin-enabled                         true
volume-type                          virtual-volume
vpd-id                               VPD83T3:6000144000000010e018b6fbc02ab396

VPlexcli:/clusters/cluster-1/virtual-volumes>

```

Example

Create a virtual volume with RAID 1 local devices and specified storage volumes:

```

VPlexcli:/> storage-tool compose --name TEST --geometry raid-1 --storage-
volumes VPD83T3:60060160cea33000fc39e04dac48e211,

```



```
VPD83T3:60060160cea33000fb9c532eac48e211,
VPD83T3:600601605a903000f2a9692fa548e211,
VPD83T3:600601605a903000f3a9692fa548e211
```

See also

- `storage-volume unclaim`
- `virtual-volume provision`

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] path
```

```
[-i|--interval] [seconds]
```

Arguments

Required arguments	
<code>[-n --director] <i>path</i></code>	* The director on which to show or change the delay for automatic unbanishment.
Optional arguments	
<code>[-i --interval] [<i>seconds</i>]</code>	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)” in the storage-volume unbanish command description.

At regular intervals, the VPLEX directors look 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 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`
- `storage-volume unbanished`

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] name
--thin-rebuild

--batch-size integer
[-d|--storage-volumes] path,path...
[-f|--force]
```

Arguments

Required arguments	
--------------------	--

<code>[-d --storage-volumes]</code> <i>path,path...</i>	* List of one or more storage volumes to claim.
Optional arguments	
<code>[--appc]</code>	Make the specified storage volumes application consistent. Prevents data already on the specified storage volumes 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 <code>set</code> 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.
<code>[-n --name]</code> <i>name</i>	The new name of the storage volume after it is claimed.
<code>--thin-rebuild</code>	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 <code>set</code> command.
<code>--batch-size</code> <i>integer</i>	When using wildcards to claim multiple volumes with one command, the maximum number of storage volumes to claim at once.
<code>[-f --force]</code>	Force the storage volume to be claimed. For use with non-interactive scripts.

* - 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).


This command can fail if there is not a sufficient number of meta volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

Thin provisioning

Thin provisioning allows storage to migrate onto a thinly provisioned storage volumes while allocating the minimal amount of thin storage container capacity.

Thinly provisioned storage volumes can be incorporated into RAID 1 mirrors with similar consumption of thin storage container 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 you to specify thin provisioning for each back-end storage volume. If a storage volume is thinly provisioned, the `thin-rebuild` attribute must be 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 container resources set aside and not be in competition for this space with any user-data volumes on the same pool.

lf:

- 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.

Examples

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	VPD83T3 ID	Capacity	Use	Vendor	ID	Status	Type	Thin
Basic_cl_ramdisk_100GB_684	VPD83T3:6000144000000103017dfea88355431	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_685	VPD83T3:6000144000000103017dfea88355433	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_686	VPD83T3:6000144000000103017dfea88355435	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_687	VPD83T3:6000144000000103017dfea88355437	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_688	VPD83T3:6000144000000103017dfea88355439	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_689	VPD83T3:6000144000000103017dfea8835543b	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_68	VPD83T3:6000144000000103017dfea88354f61	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_690	VPD83T3:6000144000000103017dfea8835543d	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_691	VPD83T3:6000144000000103017dfea8835543f	100G	claimed	EMC	alive	normal	false	false
Basic_cl_ramdisk_100GB_692	VPD83T3:6000144000000103017dfea88355441	100G	claimed	EMC	alive	normal	false	false

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> claim --
storage-volumes VPD83T3:6006016021d025007029e95b2327df11
```

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
```

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,VPD83T3:6006016091c50e005257534d0
c17e011
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,
                        0x5000144230354910/0x5006016930600523/6,
                        0x5000144230354910/0x5006016830600523/6,
                        0x5000144230354911/0x5006016830600523/6,
                        0x5000144220354910/0x5006016930600523/6,
                        0x5000144220354910/0x5006016830600523/6,
                        0x5000144220354911/0x5006016930600523/6,
                        0x5000144220354911/0x5006016830600523/6
largest-free-chunk      2G
locality                 -
operational-status      ok
storage-array-name      EMC-CLARiiON-APM00042201310
storage-volumetype      normal
system-id               VPD83T3:6006016091c50e005057534d0c17e011
thin-capable            false
thin-rebuild            false
total-free-space        2G
use                      claimed
used-by                  []
vendor-specific-name     DGC

```

Claim multiple storage volumes whose names begin with VPD83T3:600601602:

```

VPlexcli:/clusters/cluster-1> storage-volume claim --storage-volumes
VPD83T3:600601602*

```

See also

- set
- storage-volume claimingwizard
- storage-volume unclaim

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
```

```
[--force]
--appc
--thin-rebuild
```

Arguments

Optional arguments	
<code>[-c --cluster]</code> <i>cluster</i>	- Cluster on which to claim storage.
<code>[-f --file]</code> <i>file, file...</i>	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.
<code>[-d dryRun]</code>	Do a dry-run only, do not claim and name the storage volumes.
<code>[-t --set-tier]</code> <i>list</i>	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.
<code>[--force]</code>	Forces a successful run of the claimingwizard. For use with non-interactive scripts.
<code>--appc</code>	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 or extents are in the claimed state.
<code>--thin-rebuild</code>	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

You must first claim and optionally name a storage volume before using the storage volume 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’.

This command can fail if there is not a sufficient number of meta volume slots. See the troubleshooting section of the VPLEX procedures in the SolVe Desktop for a resolution to this problem.

The following table lists examples to create hint files:

Table 23 Create hints files for storage-volume naming

Storage array	Command to create hints file
Dell EMC CLARiiON	<code>navicli -h 192.168.47.27 getlun -uid -name > Clar0400.txt</code>
Dell EMC Symmetrix	<code>symdev -sid 781 list -wwn > Symm0781.txt</code>

Table 23 Create hints files for storage-volume naming (continued)

Storage array	Command to create hints file
Dell EMC VPLEX	<code>export storage-view map -f EMC_PROD12.txt -v <>\ \ views</code>
IBM DS4300	<code>SMcli 192.168.97.121 -c "show logicalDrives;" > DS4300_121.txt</code>
IBM Nextra	<code>xcli -c nextra_lab -x vol_list > Nextra_lab.txt</code>
HP EVA	<code>sssu "select manager <i>hostname</i> username=<i>username</i> password=<i>password</i>" "select system <i>systemname</i> "ls vdisk full" > <i>filename.txt</i></code>
Generic	<p>Text file of the following format:</p> <pre>> Generic storage-volumes > VPD83T3:600a0b800011ea0a000073c5468cedbd MyName1 > 600a0b800011ea0a000073c5468cedbc MyName2 > vpd83t3:600A0b800011EA0a000073c5468cEdbd MyName3</pre> <p>For generic storage volumes, names may include letters, numbers, and '_'.</p>

Example

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_llun name, and all storage volumes from the Symmetrix® array as Symm04A1H_lun name

Dell 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.

In the following example, the `claimingwizard` command with no arguments claims storage volumes from an Dell 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	VPD83T3:60000970000192602773533030354639	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:

```
Symmlast-4-digits-of-array-serial-number_Symmetrix-Device-Number
```

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.
```

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`
- `storage-volume unclaim`

storage-volume find-array

Searches storage arrays for the specified storage-volumes.

Contexts

All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is `find-array`.

Syntax

```
storage-volume find-array
[-d|--opt_s_vol] storage-volume
```

Arguments

Required arguments	
<code>[-d --opt_s_vol] storage-volume</code>	* Storage volume pattern for which to search. The pattern conforms to glob. The following pattern symbols are supported: *, ?, [<i>seq</i>], [<i>!seq</i>].

* argument is positional.

Description

Searches all the storage arrays in all clusters for the specified storage volumes.

The search is case-sensitive.

Example

Find all storage arrays for storage volumes in cluster-1:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> find-array *
Searching for cluster-1_journal
Storage-volume: cluster-1_journal is in: /clusters/cluster-1/storage-elements/
storage-arrays/EMC-Invista-14403b
Searching for cluster-1_journal_1
Storage-volume: cluster-1_journal_1 is in: /clusters/cluster-1/storage-
elements/storage-arrays/EMC-Invista-14403b
Searching for CLAR1912_10G_Aleve_1_vol_1
Storage-volume: CLAR1912_10G_Aleve_1_vol_1 is in: /clusters/cluster-1/storage-
elements/storage-arrays/EMC-CLARiiON-APM00111501912
Searching for CLAR1912_10G_Aleve_1_vol_2
Storage-volume: CLAR1912_10G_Aleve_1_vol_2 is in: /clusters/cluster-1/storage-
elements/storage-arrays/EMC-CLARiiON-APM00111501912
.
.
.
```

Find a storage array for a specified storage volume:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> find-array -d
VPD83T3:60060160d2a02c00ff3b1abb99e3e011
Searching for VPD83T3:60060160d2a02c00ff3b1abb99e3e011
Storage-volume: VPD83T3:60060160d2a02c00ff3b1abb99e3e011 is in: /clusters/
cluster-1/storage-elements/storage-arrays/EMC-CLARiiON-APM00111402062
```

Alternatively, you could enter the command as:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> find-array --
opt_s_vol VPD83T3:60060160d2a02c00ff3b1abb99e3e011
Searching for VPD83T3:60060160d2a02c00ff3b1abb99e3e011
Storage-volume: VPD83T3:60060160d2a02c00ff3b1abb99e3e011 is in: /clusters/
cluster-1/storage-elements/storage-arrays/EMC-CLARiiON-APM00111402062
```

See also

- `storage-volume claimingwizard`

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-volumes] path [,path...]
```

Arguments

Required arguments	
<code>[-d --storage-volumes] path[, path...]</code>	* List of one or more storage volumes to forget.

* - argument is positional.

Description

Storage volumes can be remembered even if a cluster is not currently in contact with them. Use this command to tell the cluster that the storage volumes are not coming back and therefore it is safe to forget about them.

You can use the `storage-volume forget` command only on storage volumes that are unclaimed or unusable, and unreachable.

This command also forgets the logical unit for this storage volume.

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 the `--verbose` argument to print a message for each volume that could not be forgotten.

Use the `logical-unit forget` command for the functionality supported by the removed arguments.

Example

Forget a specified storage volume:

```
Vplexcli:/clusters/cluster-1/storage-elements/storage-volumes> forget --
storage-volume VPD83T3:6006016021d0250027b925ff60b5de11
```

Forget all unclaimed, unused, and unreachable storage volume on the cluster:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> storage-volume
forget *
3 storage-volumes were forgotten.
```

Use the `--verbose` argument to display detailed information while you forget all unclaimed, unused, and unreachable storage volumes on the cluster:

```
VPlexcli:/clusters/cluster-1/storage-elements/storage-volumes> storage-volume
forget * --verbose
WARNING: Error forgetting storage-volume
'VPD83T3:60000970000192602773533030353933': The 'use' property of storage-
volume VPD83T3:60000970000192602773533030353933' is 'meta-data' but must be
'unclaimed' or 'unusable' before it can be forgotten.
.
.
.
3 storage-volumes were forgotten:
VPD83T3:6006016030802100e405a642ed16e111
.
.
```

See also

- `logical-unit forget`
- `storage-volume unclaim`

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] path
```

Arguments

Required arguments	
<code>[-n --director] path</code>	*The director whose banished storage volumes to display.

Description

Displays the names of storage volumes that are currently banished for a given director.

See “Banished storage volumes (LUNs)” in the `storage-volume unbanish` command description.

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
- storage-volume unbanish

storage-volume list-thin-capable

Provides a summary of all thin-capable storage-volumes and determines whether or not the volumes are declared thin (thin-rebuild).

Contexts

All contexts.

Syntax

```
storage-volume list-thin-capable
[-c|--clusters] context path [, context path...]
[-h|--help]
[--verbose]
```

Arguments

Required arguments	
<code>[-c --clusters] <i>context path</i></code>	* Specifies the clusters at which to list the thin-capable storage-volumes.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more help during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Lists all thin-capable storage volumes at the given clusters with an abbreviated list of fields for performance. The fields include: name, thin-rebuild status, capacity, current use, and I/O status. If more fields are desired, use the `--verbose` option.

Example

Displays thin-capable storage volumes for the specified clusters.

```
VPlexcli:/> storage-volume list-thin-capable --clusters cluster-1, cluster-2
cluster-1:
Name                Thin Rebuild  Capacity  Use      IO Status
-----
```

```

XtremIO0547_LUN_00010 false 10G used alive
XtremIO0547_LUN_00009 false 10G used alive
XtremIO0547_LUN_00003 false 10G used alive
thin_capable_sv_1 false 10G claimed alive
thin_capable_sv_2 false 10G used alive
XtremIO0547_LUN_00004 false 10G claimed alive
XtremIO0547_LUN_00005 false 10G claimed alive
XtremIO0547_LUN_00006 false 10G claimed alive
XtremIO0547_LUN_00007 false 10G claimed alive
XtremIO0547_LUN_00008 false 10G claimed alive

cluster-2:
Name Thin Rebuild Capacity Use IO Status
-----
VPD83T3:514f0c5d8320055e false 10G claimed alive
VPD83T3:514f0c5d83200560 false 10G claimed alive
XtremIO0541_LUN_00000 false 10G claimed alive
XtremIO0541_LUN_00002 false 10G claimed alive
XtremIO0541_LUN_00004 false 10G claimed alive
XtremIO0541_LUN_00005 false 10G claimed alive
XtremIO0541_LUN_00006 false 10G claimed alive
XtremIO0541_LUN_00007 false 10G claimed alive
XtremIO0541_LUN_00008 false 10G claimed alive
XtremIO0541_LUN_00009 false 10G claimed alive
XtremIO0541_LUN_00010 false 10G claimed alive

VPlexcli:/>

```

See also

- virtual-volume list-thin

storage-volume resurrect

Resurrect the specified storage-volumes.

Contexts

All contexts.

In `/clusters/cluster/storage-elements/storage-volumes` context, command is `resurrect`.

Syntax

```
storage-volume resurrect
[-d|--storage-volume] path [, path...]
```

```
[-f|--force]
```

Arguments

Required arguments	
<code>[-d --storage-volume] <i>path</i> [, <i>path</i>...]</code>	List of one or more storage volume with dead I/O status to resurrect.
Optional arguments	
<code>[-f --force]</code>	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.

CAUTION 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. 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 or deactivation on the storage array
- An operator shrinks the size of a storage volume, causing VPLEX to refuse to do I/O to the storage volume.
- 100% allocated thin pools
- Persistent reservation on storage volume
- Dropped frames due to a bad cable or SFP

Dead storage volumes are indicated by one of the following conditions:

- 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
```

Examples

Resurrect two storage volumes:

```
VPlexcli: /> storage-volume resurrect --storage-volumes
Symm1852_BAC, Symm1852_BA8
```

See also

- cluster status
- storage-volume forget
- storage-volume summary

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	
<code>[-c --clusters] <i>cluster</i>, [<i>cluster</i>]...</code>	Displays storage volumes for only the specified clusters.

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 clusters.

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.

Table 24 storage-volume summary field descriptions

Field	Description
Health summary (displayed only for unhealthy storage volumes)	
Name	Name of storage volume.
I/O Status	alive - I/O is proceeding normally on the storage volume.

Table 24 storage-volume summary field descriptions (continued)

Field	Description
	<p>dead - 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.</p> <p>unreachable - The storage volume is unreachable.</p>
Operational Status	<p>ok - The storage volume is functioning normally.</p> <p>degraded - 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>unknown - VPLEX cannot determine the storage volume's Operational state, or the state is invalid.</p> <p>error - VPLEX has marked the storage volume as hardware-dead.</p> <p>starting - The storage volume is not yet ready.</p> <p>lost communication - The storage volume is unreachable.</p>
Health State	<p>degraded - The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID 1 device).</p> <p>ok - The extent is functioning normally.</p> <p>non-recoverable-error - 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.</p> <p>unknown - VPLEX cannot determine the extent's Operational state, or the state is invalid.</p> <p>critical failure - VPLEX has marked the storage volume as hardware-dead.</p>
Storage-Volume Summary	
out-of-date	Of the total number of storage volumes on the cluster, the number that are out-of-date compared to their mirror.
storage-volumes	Total number of storage volumes on the cluster.

Table 24 storage-volume summary field descriptions (continued)

Field	Description
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.
Meta Slots	<p>Total - The total slots used.</p> <p>reclaimable - The number of slots that can be reclaimed.</p> <p>used - The number of slots that are used.</p> <p>storage-volume - The number of slots used for storage volumes.</p> <p>extents - The number of slots used for extents.</p> <p>logging-segments - The number of slots used for logging segments.</p>

Examples

Display default summary (all clusters) on a VPLEX with unhealthy volumes:

```

Vplexcli:/> storage-volume summary
SUMMARY (cluster-1)
StorageVolume Name           IO Status   Operational Status
Health State
-----
Clar0106_LUN14               alive       degraded
degraded
Storage-Volume Summary (no tier)
-----
Health                       out-of-date 0
                              storage-volumes 363
                              unhealthy 1
Vendor                       DGC         114
                              EMC         248
                              None        1

```

```

Use                meta-data          4
                   unusable          0
                   used              358
Capacity           total              2T
SUMMARY (cluster-2)
Storage-Volume Summary (no tier)
-----
Health             out-of-date          0
                   storage-volumes 362
                   unhealthy          0
Vendor             DGC                 114
                   EMC                 248
Use                meta-data          4
                   used              358
Capacity           total              1.99T
    
```

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
    
```

When slot usage reaches 90%, this command also displays the following:

```

Meta Slots                total          64000
                           reclaimable          9600
                           used              57600
                           storage-volumes 8000
                           extents           24000
                           logging-segments 25600
    
```

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
    
```

See also

- ds summary
- ds dd set-log
- export port summary
- export storage-view summary
- extent summary
- local-device summary
- storage-volume resurrect
- virtual-volume provision

storage-volume unbanish

Unbanishes a storage volume on one or more directors.

Contexts

In `/clusters/cluster/storage-elements/storage-volumes` context, command is unbanish.

All contexts.

Syntax

```

storage-volume unbanish
[-n|--directors] path [, path...]

[-d|--storage-volume] path
    
```

Arguments

Required arguments	
<code>[-n --directors] <i>path</i> [, <i>path</i>],...</code>	* The context path of the directors to unbanish the given storage volume on.
Optional arguments	
<code>[-d --storage-volume] <i>path</i></code>	The context path of the storage volume to unbanish.

	This argument is not required if the current context is a <code>storage-volume</code> or below. If the current context is a <code>storage-volume</code> 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 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 the system detects a LUN in this state, it waits 20 seconds for the LUN to return to normal. If the LUN does not return to the expected state, the system 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`
- `storage-volume list-banished`

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] path, [path...]
[-r|--return-to-pool]
```

Arguments

Required arguments	
<code>[-d --storage-volumes] <i>path</i>, [<i>path...</i>]</code>	* Specifies the storage volumes to unclaim.
Optional arguments	
<code>[-b --batch-size] <i>integer</i></code>	Specifies the maximum number of storage volumes to unclaim at once.
<code>[-r --return-to-pool]</code>	Returns the storage capacity of each VIAS-based volume to the pool on the corresponding storage-array.

* - argument is positional.

Description

Use the `storage-volume unclaim` command to return the specified storage volumes to the unclaimed state.

The target storage volume must not be in use.

Note: When you use the `storage-volume unclaim` command with VIAS based storage volumes, the command removes the storage volumes from VPLEX and they are no longer visible. When you use the command with non VIAS based storage volumes, the command marks the storage volumes as unclaimed. This is the intended behavior.

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	VFD83 ID	Capacity	Use	Vendor	IO Status	Type	Thin Rebuild
Basic_cl_ramdisk_100GB_684	VFD83T3:6000144000000103017dfea88355431	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_685	VFD83T3:6000144000000103017dfea88355433	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_686	VFD83T3:6000144000000103017dfea88355435	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_687	VFD83T3:6000144000000103017dfea88355437	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_688	VFD83T3:6000144000000103017dfea88355439	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_689	VFD83T3:6000144000000103017dfea8835543b	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_68	VFD83T3:6000144000000103017dfea88354f61	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_690	VFD83T3:6000144000000103017dfea8835543d	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_691	VFD83T3:6000144000000103017dfea8835543f	100G	claimed	EMC	alive	normal	false
Basic_cl_ramdisk_100GB_692	VFD83T3:6000144000000103017dfea88355441	100G	claimed	EMC	alive	normal	false

```
Vplexcli:/clusters/cluster-2/storage-elements/storage-volumes> unclaim -d
Basic_cl_ramdisk_100GB_686_See also
```

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] path [,path...]
```

Arguments

Required arguments	
<code>[-d --storage-volumes] path</code>	* 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
```

syrcollect

Collects system configuration data for System Reporting (SYR).

Contexts

All contexts.

Syntax

```
syrcollect
[-d|--directory] directory
```

Arguments

Optional arguments	
[-d --directory] <i>directory</i>	<p>Non-default directory in which to store the output. Files saved in the non-default directory are not automatically sent to Dell EMC.</p> <ul style="list-style-type: none"> • Default: Files are stored in the <code>Event_Msg_Folder</code> in the directory specified in the <code>EmaAdaptorConfig.properties</code> file. • <code>EmaAdaptorConfig.properties</code> and the <code>Event_Msg_Folder</code> are located in <code>/opt/emc/VPlex</code> on the management server. • Files in the default directory are automatically sent to Dell EMC.

Description

Manually starts a collection of SYR data, and optionally sends the resulting zip file to Dell 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 Dell EMC.

Use the `--directory` argument to specify a non-default directory. Output files sent to a non-default directory are not automatically sent to Dell EMC.

Example

Start an SYR data collection, and send the output to Dell EMC:

```
VPlexcli: /> syrcollect
```

Start an SYR data collection, and send the output to the specified directory:

```
VPlexcli: /> syrcollect -d /var/log/VPlex/cli
```

See also

- `scheduleSYR add`
- `scheduleSYR list`
- `scheduleSYR remove`

tree

Displays the context tree.

Contexts

All contexts.

Syntax

```
tree
    [-e|--expand]
    [-c|--context] subcontext-root
    [-s|--select] glob-pattern
```

Arguments

Optional arguments	
<code>[-e --expand]</code>	Expand the subcontexts.
<code>[-c --context] <i>subcontext-root</i></code>	The subcontext to use as the root for the tree.
<code>[-s --select] <i>glob-pattern</i></code>	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.

Examples

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`
- `set`

unalias

Removes a command alias.

Contexts

All contexts.

Syntax

```
unalias
  [-n|--name] name
  [-a|--all]
```

Arguments

Optional arguments	
<code>[-n --name] <i>name</i></code>	The name of the alias to remove.
<code>[-a --all]</code>	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`

user add

Adds a username to the VPLEX management server and optionally assigns a role to the added username.

Contexts

All contexts.

Syntax

```

user add
[-u|--username] username
[-r|--rolename] rolename
    
```

Arguments

Required arguments	
<code>[-u --username]</code> <i>username</i>	Username to add.

Optional arguments	
<code>[-r --rolename]</code> <i>rolename</i>	Rolename to assign.

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**.

i Note: In VPLEX Metro configurations, the system does not propagate VPLEX CLI accounts created on one management server to the second management server. The `user list` command displays only those accounts configured on the local management server, not both servers.

Administrative privileges are required to add, delete, and reset user accounts. You must reset the password for the admin account the first time you access the admin account. After the

admin password is 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 listed in the *Dell EMC VPLEX Security Configuration Guide*. A prompt to change the `admin` account password appears. Enter a new password.

Examples

Login to the CLI as an Administrator user.

At the CLI prompt, type the `user add` command:

```
Vplexcli:/> user add -u TestUser
```

A prompt for the Administrator password appears:

```
admin password:
```

Type the password for the Administrator username.

When prompted for the new password enter a new password for the username you entered:

```
New password:
```

Retype the password when prompted to confirm the password:

```
Confirm password:
```

Type the `user list` command to verify the new username is added:

```
Vplexcli:/> user list
Username
-----
admin
service
TestUser
```

See also

- `user event-server change-password`
- `user passwd`
- `user remove`
- `user reset`

user event-server add-user

Adds a user name to the event server.

Contexts

User context.

Syntax

```
user event-server add-user
    [-u|--username] username
    [-h|--help]
    [--verbose]
```

Arguments

Required arguments	
<code>[-u --username] <i>username</i></code>	* Specifies the user name to add in the event server.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Add a user to the event server to access VPLEX events.

Example

Add a user to the event server:

```
VPlexcli:/> user event-server add-user -u temp1
Enter the new password:
Confirm new password:
The username "temp1" was successfully added in the event-server.
```

See also

- `user event-server change-password`

user event-server change-password

Changes the password for a user accessing the event server.

Contexts

User context.

Syntax

```
user event-server change-password
    [-u|--username] username
    [-h|--help]
    [--verbose]
```

Arguments

Required arguments	
--------------------	--

<code>[-u --username] <i>username</i></code>	* Specifies the user name for which to change the password.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Change the password that an external subscriber uses to access VPLEX events.

Examples

An event-server user changes the default password:

```
VPlexcli:/> user event-server change-password -u temp1
Enter current password:
Enter the new password:
Confirm new password:
```

See also

- `user event-server add-user`

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 configurations, the system does not propagate VPLEX CLI accounts created on one management server to the second management server. The `user list` command displays only accounts configured on the local management server, not on both servers.

Examples

Display the user accounts configured on the local management server:

```
VPlexcli:/> user list
Username
-----
admin
service
TestUser
```

See also

- user add
- user passwd
- user remove
- user reset

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	
<code>[-u --username] <i>username</i></code>	*Username for which to change the password.

* - argument is positional.

Description

Executable by all users to change the password only for their own username.

Examples

```
VPlexcli: /> user passwd -u TestUser
```

```
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`
- `user event-server change-password`
- `user remove`
- `user reset`

user remove

Removes a username from the VPLEX management server.

Contexts

All contexts.

Syntax

```
user remove
    [-u|--username] username
```

Arguments

Required arguments	
<code>[-u --username] username</code>	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 listed in the *Dell EMC VPLEX Security Configuration Guide*. 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 -u TestUser
admin password:
```

Type the password for the Administrator username.

The specified user username is removed.

Type the `user event-server change-password` command to verify the removal.

See also

- `user add`
- `user event-server change-password`
- `user passwd`
- `user reset`

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	
<code>[-u --username] <i>username</i></code>	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 listed in the *Dell EMC VPLEX Security Configuration Guide*. 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.

Examples

Login as an Administrator user.

Type the `user reset --username username` command:

```
VPlxcli:/> user reset --username TestUser
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`
- `user event-server change-password`
- `user passwd`
- `user remove`

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.

Examples

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 -d' for details.
Validate meta-volume
Checking cluster cluster-1 ...
Checking cluster cluster-2 ...
ok
```

See also

- cluster status
- connectivity validate-be
- health-check

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

<code>[-c --cluster]</code> cluster	Specify the cluster on which to start cache vaulting.
<code>[--force]</code>	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.

Examples

Start a manual vault on cluster-1:

```
Vplexcli: /> vault go --cluster cluster-1
```

See also

- `vault overrideUnvaultQuorum`
- `vault status`
- *Dell EMC VPLEX Administration Guide*

vault overrideUnvaultQuorum

Allows the cluster to proceed with the recovery of the vaults without all the required directors.

Contexts

All contexts.

Syntax

```
vault overrideUnvaultQuorum
    [-c|--cluster] cluster
    [--evaluate-override-before-execution]
    [--force]
```

Arguments

<code>[-c --cluster] cluster</code>	Overrides unvault quorum for the specified cluster.
<code>--evaluate-override-before-execution</code>	Evaluates the possible outcome of running this command but does not do anything.
<code>--force</code>	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 directors 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.

Examples

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
```

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!
```

In the following example, the command evaluates the cluster vault status and overrides the unvault quorum when:

- 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!
```

- 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 go`
- `vault status`
- *Dell EMC VPLEX Administration Guide*
- VPLEX procedures in the SolVe Desktop

vault status

Displays the current cache vault/unvault status of the cluster.

Contexts

All contexts.

Syntax

```
vault status
    [-c|--cluster] cluster
    [--verbose]
```

Arguments

Optional arguments	
<code>[-c --cluster] <i>cluster</i></code>	Displays vault status for the specified cluster.
<code> [--verbose]</code>	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 25 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 directors 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.

Examples

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
```

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
```

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
```

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
- vault overrideUnvaultQuorum

- *Dell EMC VPLEX Administration 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

Description

Verifies that the Fibre Channel switch on each cluster's internal management network has been configured correctly.

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...
    [--verbose]
```

Arguments

Optional arguments	
[-a --all]	Displays version information for all connected directors.

<code>[-n --directors] <i>context-path</i></code>	* Display version information for only the specified directors.
<code>--verbose</code>	Displays version information for individual software components on each director.

* - argument is positional.

Description

This command displays version information for all directors, a specified director, or individual software components for each director.

Table 26 Software components

Component Abbreviation	Description
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 Dell EMC Common Object Manager is a hub of communications and common services for applications based on Dell 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 clusters.
FW Bundle Rev	Firmware revision.
POST Rev	Power On Self Test revision.
BIOS Rev	Boot firmware revision.

Table 26 Software components (continued)

Component Abbreviation	Description
SSD Model	Solid state disk drive model information.

Examples

Display management server/SMS version information:

```

VPLEXcli:/> version
What                Version          Info
-----
Product Version    5.4.0.00.00.10 -
SMSv2              D35.20.0.10.0  -
Mgmt Server Base   D35.20.0.1     -
Mgmt Server Software D35.20.0.13    -

```

Display management server/SMS version and version for the specified director:

```

VPLEXcli:/> version director-2-1-B
What                Version          Info
-----
Product Version    5.4.0.00.00.10 -
SMSv2              D35.20.0.10.0  -
Mgmt Server Base   D35.20.0.1     -
Mgmt Server Software D35.20.0.13    -
/engines/engine-2-1/directors/director-2-1-B 6.5.54.0.0     -

```

Display version information for management server, SMS, and all directors:

```

VPLEXcli:/> version -a
What                Version          Info
-----
Product Version    5.4.0.00.00.10 -
SMSv2              D35.20.0.10.0  -
Mgmt Server Base   D35.20.0.1     -
Mgmt Server Software D35.20.0.13    -
/engines/engine-2-1/directors/director-2-1-B 6.5.54.0.0     -
/engines/engine-2-1/directors/director-2-1-A 6.5.54.0.0     -
/engines/engine-1-1/directors/director-1-1-B 6.5.54.0.0     -
/engines/engine-1-1/directors/director-1-1-A 6.5.54.0.0     -

```

Display version information for individual software components on each director. See Software components table below for a description of the components.

```

VPLEXcli:/> version -a --verbose
Product Version:  5.4.0.00.00.10
What:             SMSv2
Version:          D35.20.0.10.0
Build time:       June 09, 2014 at 11:38:36PM EDT
Build machine:    dudleyed05
Build OS:         Linux version 2.6.27-7-generic on amd64
Build compiler:   1.6.0_45
Build source:     /spgear/spgear_misc/htdocs/harness/release/1795/work/ui/src
What:             Mgmt Server Base
Version:          D35.20.0.1
What:             Mgmt Server Software
Version:          D35.20.0.13

```

```

For director /engines/engine-2-1/directors/director-2-1-B:
What: O/S
Version: D35.20.0.1 (SLES11)
What: Director Software
Version: 6.5.54.0.0
What: ECOM
Version: 6.5.1.0.0-0
What: VPLEX Splitter
Version: 4.1.b_vplex_D35_00_Ottawa_MR1.10-1
What: ZECL
Version: 6.5.52.0.0-0
What: ZPEM
Version: 6.5.52.0.0-0
What: NSFW
Version: 65.1.54.0-0
What: BIOS Rev
Version: 08.50
What: POST Rev
Version: 43.80
What: FW Bundle Rev
Version: 12.60
What: SSD Model: P30056-MTFDBAA056SAL 118032803
Version: 0005
For director /engines/engine-2-1/directors/director-2-1-A:
What: O/S
Version: D35.20.0.1 (SLES11)
What: Director Software
Version: 6.5.54.0.0
What: ECOM
Version: 6.5.1.0.0-0
What: VPLEX Splitter
Version: 4.1.b_vplex_D35_00_Ottawa_MR1.10-1
What: ZECL
Version: 6.5.52.0.0-0
What: ZPEM
Version: 6.5.52.0.0-0
What: NSFW
Version: 65.1.54.0-0
What: BIOS Rev
Version: 08.50
What: POST Rev
Version: 43.80
What: FW Bundle Rev
Version: 12.60
What: SSD Model: P30056-MTFDBAA056SAL 118032803
Version: 0005
For director /engines/engine-1-1/directors/director-1-1-B:
What: O/S
Version: D35.20.0.1 (SLES11)
What: Director Software
Version: 6.5.54.0.0
What: ECOM
Version: 6.5.1.0.0-0
What: VPLEX Splitter
Version: 4.1.b_vplex_D35_00_Ottawa_MR1.10-1
What: ZECL
Version: 6.5.52.0.0-0
What: ZPEM
Version: 6.5.52.0.0-0
What: NSFW
Version: 65.1.54.0-0
What: BIOS Rev
Version: 08.50
What: POST Rev
Version: 43.80
What: FW Bundle Rev
Version: 12.60
What: SSD Model: P30056-MTFDBAA056SAL 118032803

```

```

Version: 0005
For director /engines/engine-1-1/directors/director-1-1-A:
What: O/S
Version: D35.20.0.1 (SLES11)
What: Director Software
Version: 6.5.54.0.0
What: ECOM
Version: 6.5.1.0.0-0
What: VPLEX Splitter
Version: 4.1.b_vplex_D35_00_Ottawa_MR1.10-1
What: ZECL
Version: 6.5.52.0.0-0
What: ZPEM
Version: 6.5.52.0.0-0
What: NSFW
Version: 65.1.54.0-0
What: BIOS Rev
Version: 08.50
What: POST Rev
Version: 43.80
What: FW Bundle Rev
Version: 12.60
What: SSD Model: P30056-MTFDBAA056SAL 118032803
Version: 0005

```

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
[-n | --thin]
[-i | --initialize]
[--confirm-init]
[--verbose]

```

Arguments

Required arguments	
<code>[-r --device] <i>context-path</i></code>	* Device on which to host the virtual volume.
Optional arguments	
<code>[-t --set-tier] <i>tier</i></code>	Set the storage-tier for the new virtual volume.
<code>[-n --thin]</code>	Specifies whether to create a thin-enabled virtual volume or not.
<code>[-i --initialize]</code>	Initializes the virtual volume by erasing 10 MB of the initial storage blocks. This prevents the virtual volume from retaining old or stale data. This must be used along with the <code>confirm-init</code> option.

<code>[--confirm-init]</code>	Confirms the initialization process on the virtual volume. This must be used along with the <code>initialize</code> option.
-------------------------------	---

* - 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.

Table 27 virtual-volume field descriptions

Field	Description
<code>block count</code>	The number of blocks in the volume.
<code>block size</code>	The size of a single block, in kilobytes.
<code>cache-mode</code>	Synchronous (write-through).
<code>capacity</code>	The total number of bytes in the volume. Equals the <code>block-size</code> multiplied by the <code>block-count</code> . Note: The capacity of a virtual volume on which the initialization process has failed will be 0.
<code>thin-capable</code>	Determines whether the virtual volume is thin-capable or not.
<code>thin-enabled</code>	Determines whether the virtual volume is configured as thin-enabled.
<code>consistency-group</code>	The name of the consistency group to which this volume belongs, if any.
<code>expandable-capacity</code>	Excess capacity not yet exposed to the host by the virtual volume. This capacity is available for expanding the virtual volume. <ul style="list-style-type: none"> Zero (0) - Expansion is not supported on the virtual volume or that there is no capacity available for expansion. Non-zero - The capacity available for virtual volume expansion using the <code>storage-volume</code> method.
<code>expansion-method</code>	The expansion method that can be used to expand the virtual volume.

Table 27 virtual-volume field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>concatenation</code> - The virtual volume can be expanded only by adding the specified extents. • <code>not-supported</code> - The virtual volume cannot be expanded. This includes virtual volumes that are members of RecoverPoint consistency groups. • <code>storage-volume</code> - The virtual volume can be expanded using storage array based volume expansion or by migrating to a larger device.
<code>expansion-status</code>	<p>Expansion status for the volume.</p> <ul style="list-style-type: none"> • <code>-</code> - None of the other expansion states apply. No operation is blocked by this state. • <code>failed</code> - An expansion has failed. The expansion has failed and the expansion must be re-tried. If the expansion is not retried this state will persist for up-to 2 days. See <code>health-indications</code> for more information. • When an expansion fails, the overall health, <code>operational-status</code>, or <code>service-status</code> of the virtual-volume is not degraded. • <code>in-progress</code> - An expansion has been started, but has not completed. The following operations are blocked on the volume: additional expansion, migration, and NDU. • <code>unknown</code> - VPLEX could not determine the expansion status of the volume.
<code>health-indications</code>	<p>Indicates the reasons for:</p> <ul style="list-style-type: none"> • A health-state that is not 'ok' • The reasons for the failure of virtual volume expansion or initialization.
<code>health state</code>	<ul style="list-style-type: none"> • <code>major failure</code> - One or more of the virtual volume's underlying devices is out-of-date, but will never rebuild. • <code>minor failure</code> - One or more of the virtual volume's underlying devices is out-of-date, but will rebuild.

Table 27 virtual-volume field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>non-recoverable error</code> - VPLEX cannot determine the virtual volume's Health state. • <code>ok</code> - The virtual volume is functioning normally. • <code>unknown</code> -VPLEX cannot determine the virtual volume's Health state, or the state is invalid.
<p><code>initialization-status</code></p>	<p>Status of the initialization process on the virtual volume.</p> <ul style="list-style-type: none"> • <code>success</code> - Indicates that the initialization process is completed successfully. • <code>failed</code> - Indicates that the initialization process is failed. • <code>in-progress</code> - Indicates that the initialization process is in progress. • <code>unknown</code> - Indicates that the initialization process remains in a status other than <code>success</code>, <code>failed</code>, or <code>in-progress</code>. <p>Note: If initialization is requested during the creation of the virtual volume, you must wait until the initialization process is completed successfully to use the virtual volume. If the initialization process fails, restart the process by using the <code>virtual-volume re-initialize</code> command.</p>
<p><code>locality</code></p>	<ul style="list-style-type: none"> • <code>local</code> - The virtual volume relies completely on storage at its containing cluster. • <code>remote</code> - 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. • <code>distributed</code> - 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.

Table 27 virtual-volume field descriptions (continued)

Field	Description
operational status	<ul style="list-style-type: none"> • <code>degraded</code> - The virtual volume may have one or more out-of-date devices that will eventually rebuild. • <code>error</code> - One or more of the virtual volume's underlying devices is hardware-dead. • <code>ok</code> - The virtual volume is functioning normally. • <code>starting</code> -The virtual volume is not yet ready. • <code>stressed</code> - One or more of the virtual volume's underlying devices is out-of-date and will never rebuild. • <code>unknown</code> - VPLEX cannot determine the virtual volume's Operational state, or the state is invalid.
recoverpoint-protection-at	The list of VPLEX cluster names at which virtual volumes are protected by RecoverPoint. For MetroPoint protected VPLEX volumes, this field contains the names of both VPLEX clusters, indicating protection at both sites.
recoverpoint-usage	<p>The replication role this virtual volume is being used by attached RecoverPoint clusters, if any.</p> <ul style="list-style-type: none"> • <code>Production Source</code> - Volumes that are written to by the host applications. Writes to production volumes are split such that they are sent to both the normally designated volumes and RPAs. • <code>Local Replica</code> - Local volumes to which production volumes replicate. • <code>Remote Replica</code> - remote volumes to which production volumes replicate. • <code>Journal</code> - Volumes that contain data waiting to be distributed to target replica volumes and copies of the data previously distributed to the target volumes. • <code>Repository</code> - A volume dedicated to RecoverPoint for each RPA cluster that stores configuration information about the RPAs and RecoverPoint consistency groups.

Table 27 virtual-volume field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> - The volume is not currently being used by RecoverPoint.
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	<p>The service status of a virtual-volume.</p> <ul style="list-style-type: none"> <code>running</code> - I/O is running for the virtual-volume. <code>inactive</code> - The virtual-volume is part of an inactive storage-view and is not visible from the host. <code>unexported</code> - The virtual-volume is unexported. <code>suspended</code> - I/O is suspended for the virtual-volume. <code>cluster-unreachable</code> - Cluster is unreachable at this time. <code>need-resume</code> - Issue re-attach to resume after link has returned. <code>disconnected</code> - Applies to distributed volumes used with RecoverPoint as replica volumes, or in production failure scenarios. It will clear when the volume is no longer used as a RecoverPoint replica, or after production failback is complete.
storage-array-family	The family of the storage array from which the virtual volume was created.
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.
vpd-id	The VPD identifier for the virtual volume.

About storage tier IDs

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.

Examples

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/r0_C1_VATS_00001_vol

VPlexcli:/clusters/cluster-1/virtual-volumes> cd r0_C1_VATS_00001_vol
VPlexcli:/clusters/cluster-1/virtual-volumes/r0_C1_VATS_00001_vol> ll
Name                               Value
-----
block-count                         20971520
block-size                          4K
cache-mode                          synchronous
capacity                             80G
consistency-group                   -
expandable                          true
expandable-capacity                 0B
expansion-method                    storage-volume
expansion-status                    -
health-indications                  []
health-state                        ok
initialization-status               success
locality                            local
operational-status                  ok
recoverpoint-protection-at          []
recoverpoint-usage                  -
scsi-release-delay                  0
service-status                      unexported
storage-array-family                clariion
storage-tier                        -
supporting-device                   r0_C1_VATS_00001
system-id                           r0_C1_VATS_00001_vol
thin-capable                        false
thin-enabled                        unavailable
volume-type                         virtual-volume
vpd-id                              VPD83T3:6000144000000010200ecb6260b7ac42

VPlexcli:/clusters/cluster-1/virtual-volumes/r0_C1_VATS_00001_vol>

```

See Also

- `virtual-volume destroy`
- `virtual-volume expand`
- `virtual-volume provision`
- `virtual-volume reinitialize`

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	
<code>[-v --virtual-volumes]</code> <i>context-path, context-path...</i>	List of one or more virtual volumes to destroy. Entries must be separated by commas. The specified virtual volumes must not be exported to hosts.
Optional arguments	
<code>[-f --force]</code>	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.

Examples

```
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`
- `virtual-volume expand`
- `virtual-volume provision`

virtual-volume expand

Non-disruptively increases the capacity of an existing virtual volume.

Contexts

All contexts.

In `clusters/cluster/virtual-volumes/ context` and below, command is `expand`.

Syntax

```
virtual-volume expand
  [-v|--virtual-volume] context-path
  [-e|--extent] extent
  [-f|--force]
```

Arguments

Required arguments	
<code>[-v --virtual-volume] <i>context-path</i></code>	<p>* The virtual volume to expand.</p> <ul style="list-style-type: none"> For both storage volume and concatenation methods of expansion, the virtual volume must be expandable, and have a geometry of RAID 1, RAID C, or RAID 0. For storage-volume expansions, the virtual volume must be expandable, and have a geometry of RAID 1, RAID C, RAID 0, or DR1.
Optional arguments	
<code>[-e --extent] <i>extent</i></code>	<p>* The target local device or extent to add to the virtual volume using the concatenation method of expansion. The local device or extent must not have a virtual volume on top of it.</p>
<code>[-f --force]</code>	<p>The meaning of this argument varies, depending on whether the <code>--extent</code> argument is used (expansion method = concatenation) or not used (expansion-method = storage-volume)</p> <ul style="list-style-type: none"> For storage-volume expansion, the <code>--force</code> argument skips the confirmation message. For concatenation expansion, the <code>--force</code> argument expands a virtual volume built on a RAID 1 device using a target that is not a RAID 1 or that is not as redundant as the device supporting the virtual volume.

* - argument is positional.

Description

This command expands the specified virtual volume using one of two methods; storage-volume or concatenation.

The `ll` command output shows whether the volume is expandable, the expandable capacity (if any), and the expansion method available for the volume. For example:

```
VPlexcli:> ll /clusters/cluster-1/virtual-volumes/ Test_volume
Name                               Value
-----
.
.
.
capacity                           0.5G
consistency-group                   -
expandable                          true
expandable-capacity                 4.5G
expansion-method                    storage-volume
expansion-status                     -
.
.
.
```

There are two methods to expand a virtual volume; storage-volume and concatenation.

- **storage-volume** - If the virtual volume has a non-zero expandable-capacity, this command will expand the capacity of the virtual volume by it's full expandable-capacity. To use the storage-volume method of expansion, use this command without the `--extent` argument. The storage-volume method of expansion adds the entire amount of the expandable-capacity to the volume's configured capacity.
- **concatenation** - (also known as RAID C expansion) Expand the virtual volume by adding the specified extents or devices. The concatenation method does not support non-disruptive expansion of DR1 devices. Use this command with the `--extent` argument to expand a virtual volume using the concatenation method of expansion.

Note: You cannot expand a virtual volume if the initialization status of the virtual volume is failed or in-progress.

Before expanding a storage volume, understand the limitations of the function and the prerequisites required for volumes to be expanded. See the *Dell EMC VPLEX Administration Guide* for more information on how expansion works. For procedure to expand virtual volumes, see the VPLEX procedures in the SolVe Desktop.

Examples

Expand a volume using the storage-volume method:

- The `ll clusters/cluster-1/virtual-volumes` command displays virtual volumes, and whether the volumes are expandable, and the expandable capacity, if any (not all columns are shown in example).
- The `ll clusters/cluster-1/virtual-volumes/virtual-volume` command displays the method (storage-volume) of expansion applicable to the volume.
- The `expand` command starts the expansion of the specified virtual volume.
- The `ll clusters/cluster-1/virtual-volumes` command displays the expanded volume:

```
VPlexcli:/clusters/cluster-1/virtual-volumes> ll /clusters/cluster-1/
virtual-volumes
/clusters/cluster-1/virtual-volumes:
Name                               ...Capacity Locality Supporting Cache
Expandable Expandable ...
...
Device Mode
```

```

                Capacity  ...
-----
...
Raid0_1Ga_11_vol  ...5G      local   raid1-dev  synchronous
true              4.5G
RaidC_1Gb_11_vol  ...5G      local   raid1-dev  synchronous
true              0B
Test_volume       ...0.5G     local   Test       synchronous
true              4.5G
.
.
.
Vplexcli:/clusters/cluster-1/virtual-volumes> ll /clusters/cluster-1/
virtual-volumes/ Test_volume
Name              Value
-----
block-count       131072
block-size        4K
cache-mode        synchronous
capacity          0.5G
consistency-group -
expandable        true
expandable-capacity 4.5G
expansion-method  storage-volume
expansion-status  -
.
.
.
Vplexcli:/clusters/cluster-1/virtual-volumes> expand -v Test_volume/
Virtual Volume expansion can take some time and once started, cannot be
cancelled. Some operations such as upgrades and data migrations will not
be possible during the expansion. In some cases hosts and their
applications may need to be restarted once the expansion has completed.
Do you wish to proceed ? (Yes/No) yes
The expansion of virtual-volume 'Test_volume' has started.
Vplexcli:/clusters/cluster-1/virtual-volumes> cd Test_volume/
Vplexcli:/clusters/cluster-1/virtual-volumes/Test_volume> ll
Name              Value
-----
block-count       131072
block-size        4K
cache-mode        synchronous
capacity          0.5G
consistency-group -
expandable        true
expandable-capacity 4.5G
expansion-method  storage-volume
expansion-status  in-progress
health-indications []
.
.
.
Vplexcli:/clusters/cluster-1/virtual-volumes> ll /clusters/cluster-1/
virtual-volumes
/clusters/cluster-1/virtual-volumes:
Name              ...Capacity Locality Supporting Cache
Expandable Expandable ...
...
                Capacity  ...
-----
...
Raid0_1Ga_11_vol  ...5G      local   raid1-dev  synchronous
true              4.5G
RaidC_1Gb_11_vol  ...5G      local   raid1-dev  synchronous
true              0B
Test_volume       ...5G      local   Test       synchronous
true              0B
.
.
.

```

Expand a virtual volume using the concatenation method:

- The `ll clusters/cluster-1/virtual-volumes` command displays the available virtual volumes, and whether the volumes are expandable.
- The `ll clusters/cluster-1/virtual-volumes/volume-name` command displays that the concatenation method must be used to expand the target volume.
- The `ll /clusters/cluster-1/storage-elements/extents` command displays available extents.
- 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  Health   ... .. Expandable
-----                               Status      State   ... .. -----
-----                               -----   -----   ... .. -----
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/virtual-volumes/Test-Device_vol
Name                               Value
-----                               -----
.
.
.
expandable                          true
expansion-method                     concatenation
health-indications                   []
.
.
.
VPlexcli:/> ll /clusters/cluster-1/storage-elements/extents
/clusters/cluster-1/storage-elements/extents:
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

- `batch-migrate pause`
- `batch-migrate resume`
- `dm migration pause`
- `dm migration resume`
- `virtual-volume create`

- `virtual-volume destroy`
- *Dell EMC VPLEX Administration Guide*

virtual-volume list-thin

Lists the virtual volumes at the given clusters with additional thin-property filtering options.

Contexts

All contexts.

Syntax

```
virtual-volume list-thin
-t | --clusters context path
-e | --enabled true|false
-c | --capable true|false
[--verbose]
```

Arguments

Required arguments	
<code>-t --clusters <i>context path</i></code>	* The target cluster where virtual volumes are listed.
Optional arguments	
<code>-e --enabled <i>true false</i></code>	Filters volumes with the matching thin-enabled value. The value can be true or false. If omitted, the results will match volumes regardless of whether they are thin-enabled or not.
<code>-c --capable <i>true false</i></code>	Filters volumes with the matching thin-capable value. The value can be true or false. If omitted, the results will match volumes regardless of whether they are thin-capable or not.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more help during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

This command lists virtual volumes at the given clusters with additional thin-property filtering options.

The following table describes the filter combinations, and the results that are listed.

Thin-capable	Thin-enabled	Results
True	true	Volumes that are both thin-capable and thin-enabled.
False	unavailable	Volumes that are thin-enabled, but not thin-capable. This is true only for thin-to-thin or thin-to-thick migrations.

Thin-capable	Thin-enabled	Results
True	disabled	Volumes that are thin-capable, but not thin-enabled.
False	unavailable	Thick volumes. Neither thin-capable nor thin-enabled.
Unspecified?	unavailable	All volumes that are not thin-enabled.
False	unavailable	All volumes that are not thin-capable (enabled or not).

See also

`storage-volume list-thin-capable`

virtual-volume provision

Provisions new virtual volumes using storage pools.



Contexts

All contexts.

Syntax

```
virtual-volume provision
    [-b|--base-name]      base-name
    [-p|--storage-pools] storage-pools [,storage-pools...]
    [-n|--number-of-volumes] number-of-volumes
    [-c|--capacity]      capacity
    [-j|--job]            =job
    [-g|--consistency-group] consistency-group
    [-v|--storage-views] storage-views [,storage-views...]
    [-t|--thin]
    [-h|--help]
    [--verbose]
```

Arguments

Required arguments	
<code>[-b --base-name]</code> base-name	* Specifies the base-name of the virtual-volumes to provision. Refer to the Dell EMC Simple Support Matrix for a list of supported arrays.
<code>[-p --storage-pools]</code> <i>storage-pools</i> [, <i>storage-pools</i>]	* Specifies the storage-pools to use for provisioning. Storage pools must be given in <i>cluster-name</i> , <i>storage-array-name</i> , <i>storage-pool-name</i> [, <i>cluster-name</i> , <i>storage-array-name</i> , <i>storage-pool-name</i> ...] tuple format. A maximum of four tuples can be specified.  Note: A <i>storage-group</i> name is mandatory for VMAX3 arrays  Note: Storage pool names that contain spaces must be enclosed by double quotation marks.
<code>[-n --number-of-volumes]</code> <i>number-of-volumes</i>	* Specifies the number of volumes to provision. Each volume will have the same RAID geometry and capacity. A maximum of 10 volumes can be provisioned at one invocation.

<code>[-c --capacity]</code> <i>capacity</i>	* Specifies the capacity of the virtual-volumes in MB, GB, TB, and so on. The minimum size is 150 MB.
Optional arguments	
<code>[-j --job] =<i>job</i></code>	The name of the job to create.
<code>[-g --consistency-group] <i>consistency-group</i></code>	Specifies the context path of the consistency-group to add the virtual-volumes to. Note: During Integrated Array Services (VIAS) virtual volume provisioning, the virtual volume can be added to new or existing VPLEX consistency groups using the GUI. However, the CLI only supports adding virtual volumes to existing VPLEX consistency groups.
<code>[-v --storage-views]</code> <i>storage-views</i> [, <i>storage-views</i>]	Specifies the storage-views to which to add the virtual-volumes.
<code>[-t --thin]</code>	Provisions thin storage volumes on the target array and then creates thin virtual volumes using the newly provisioned thin storage volumes. The target array must have storage pools capable of provisioning thin storage volumes. For RAID 1 virtual volumes to be thin enabled all mirror legs must be from thin storage volumes. For virtual volumes created from XtremIO storage volumes, the thin-enabled property is set to true.
<code>[-h --help]</code>	Displays command line help.
<code>[--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

Provisions new virtual volumes using storage pools, and optionally specify a job name and which VPLEX storage-group to use on the back-end managed-array . The RAID geometry is determined automatically by the number and location of the specified pools.

Provisions new virtual-volumes using storage-pool *tuples* where an optional fourth *storage-group-name* value may be provided in the tuple. This would be the storage-group of masking view on the - array which exposes volumes to this cluster.

If no *storage-group* is specified, the system finds all masking views on the back-end managed arrays which expose volumes/LUNs to the specified clusters, and uses the masking view with the lowest LUN count.

Note: When creating a storage group using VIAS, always associate it with a pool. Failure to associate a storage group to a pool could cause the storage group to not work properly.

See also

- `storage-tool compose`
- `storage-volume unclaim`

virtual-volume re-initialize

Restarts the initialization process on a virtual volume.

Contexts

All contexts.

Syntax

```
virtual-volume re-initialize
    [-v | --virtual-volume] virtual-volume
    [--verbose]
```

Arguments

Required arguments	
<code>[-v --virtual-volume] <i>virtual-volume</i></code>	* The virtual-volume that you want to reinitialize.
Optional arguments	
<code>[-h --help]</code>	Displays command line help.
<code> [--verbose]</code>	Provides more output during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

This command restarts a failed initialization process on a virtual-volume. The command runs only if the `initialization-status` field of the virtual volume shows `failed`.

See Also

- `virtual-volume create`
- `virtual-volume destroy`
- `virtual-volume expand`
- `virtual-volume provision`

virtual-volume set-thin-enabled

Sets the thin-enabled property to either true or false for the given virtual volumes.

Contexts

All contexts.

Syntax

```
virtual-volume set-thin-enabled
-v | --virtual-volumes context path [, context path...]
-t | --thin-enabled arg
[-h|--help]
```

[--verbose]

Arguments

Required arguments	
-t --thin-enabled <i>arg</i>	Specifies the desired value of the thin-enabled property.
-v --virtual-volumes <i>context path</i> [, <i>context path...</i>]	* Specifies the virtual volumes for which the thin-enabled property must be set.
Optional arguments	
[-h --help]	Displays command line help.
[--verbose]	Provides more help during command execution. This may not have any effect for some commands.

* - argument is positional.

Description

This command sets the thin-enabled property to either true or false for the given virtual volumes. Virtual volumes can be specified as a parameter, using globbing or wildcards.

The `virtual-volume set-thin-enabled` command does not fail even if virtual volumes are not thin-capable. Virtual volumes that are not thin-capable are skipped. For brevity of the user messages, the regular output of this command only includes:

- the number of volumes that are set as thin-enabled (or not set)
- the number of volumes that are skipped

If you want detailed output showing exactly which volumes are set as thin-enabled or skipped, use the `--verbose` option. However, the output can be very long.

Example

Displays all the virtual volumes that are set as thin-enabled, or are skipped.

```
VPlexcli:/> virtual-volume set-thin-enabled true --virtual-volumes /clusters/
cluster-1/** --verbose

Virtual-volumes that were set thin-enabled:
thin_vol_1, thin_vol_2

Virtual-volumes that were skipped because they are either already thin-
enabled or not thin-capable:
thick_vol_1, thick_vol_2

VPlexcli:/>
```

See also

`storage-volume list-thin-capable`

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

Optional arguments	
<code>[-c --clusters] cluster</code>	List of one or more names of clusters. Display information for only the specified clusters. 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`, and total capacity for the cluster.

Displays any volumes with an expandable capacity greater than 0, and whether an expansion is in progress.

If the `--clusters` argument is not specified and the command is executed at or below a `/clusters/cluster` context, information is displayed for the current cluster.

Otherwise, virtual volumes of all clusters are summarized.

Table 28 virtual-volume summary field descriptions

Field	Description
Virtual-volume health summary (displayed only for unhealthy volumes)	
<code>volume name</code>	Name of the virtual volume.
<code>health state</code>	<ul style="list-style-type: none"> <code>major failure</code> - One or more of the virtual volume's underlying devices is out-of-date, but will never rebuild. <code>minor failure</code> - One or more of the virtual volume's underlying devices is out-of-date, but will rebuild. <code>non-recoverable error</code> - VPLEX cannot determine the virtual volume's Health state. <code>ok</code> - The virtual volume is functioning normally. <code>unknown</code> - VPLEX cannot determine the virtual volume's Health state, or the state is invalid.
<code>operational status</code>	<ul style="list-style-type: none"> <code>degraded</code> - The virtual volume may have one or more out-of-date devices that will eventually rebuild.

Table 28 virtual-volume summary field descriptions (continued)

Field	Description
	<ul style="list-style-type: none"> • <code>error</code> - One or more of the virtual volume's underlying devices is hardware-dead. • <code>ok</code> - The virtual volume is functioning normally. • <code>starting</code> -The virtual volume is not yet ready. • <code>stressed</code> - One or more of the virtual volume's underlying devices is out-of-date and will never rebuild. • <code>unknown</code> - VPLEX cannot determine the virtual volume's Operational state, or the state is invalid.
<p><code>service status</code></p>	<p>The service status of a virtual-volume.</p> <ul style="list-style-type: none"> • <code>running</code> - I/O is running for the virtual-volume. • <code>inactive</code> - The virtual-volume is part of an inactive storage-view and is not visible from the host. • <code>unexported</code> The virtual-volume is unexported. • <code>suspended</code> - I/O is suspended for the virtual-volume. • <code>cluster-unreachable</code> - Cluster is unreachable at this time. • <code>need-resume</code> - Issue re-attach to resume after link has returned. • <code>disconnected</code> - Applies to distributed volumes used with RecoverPoint as replica volumes, or in production failure scenarios. It will clear when the volume is no longer used as a RecoverPoint replica, or after production failback is complete.
<p>Summaries</p>	
<p>Total</p>	<p>Total number of virtual volumes on the cluster, and number of unhealthy virtual volumes.</p>
<p>Locality summary</p>	<p><code>distributed</code> - Number of distributed virtual volumes.</p> <ul style="list-style-type: none"> • <code>local</code> - Number of local virtual volumes. • <code>remote</code> - Number of remote volumes.

Table 28 virtual-volume summary field descriptions (continued)

Field	Description
Cache-mode summary	synchronous - number of virtual volumes with synchronous cache mode.
Expansion summary	<ul style="list-style-type: none"> virtual-volume name - Name of any volume with expandable capacity greater than 0 or an expansion underway. expandable-capacity - Additional capacity (if any) added to the back end storage volume not yet added to the VPLEX virtual volume. capacity - Current capacity of the virtual volume. expansion-status - Indicates whether an expansion is possible is in progress, or has failed. A value of "-" indicates expansion is possible, but is not in progress, and has not failed.

Examples

In the following example, all devices on cluster-1 are healthy:

```

Vplexcli:/clusters/cluster-1/virtual-volumes> virtual-volume summary
Virtual-volume health summary (cluster-1):
    Total 589 virtual-volumes, 0 unhealthy.
Expansion summary:
    No expansion activity.
Locality summary:
    distributed : 65 virtual-volumes.
    local      : 494 virtual-volumes.
    remote    : 30 virtual-volumes.
Cache-mode summary:
    asynchronous : 0 virtual-volumes.
    synchronous  : 589 virtual-volumes.
Total virtual-volume capacity is 87.9T.
    
```

In the following example, one distributed virtual volume has expandable capacity at both clusters:

```

Vplexcli:/> virtual-volume summary
Virtual-volume health summary (cluster-1):
    Total 2152 virtual-volumes, 0 unhealthy.
Expansion summary:
virtual-volume name  expandable-capacity  capacity  expansion-status
-----
dr_one2one_CX_0_vol  10G                10G      -
Total 1 expansion: 0 in-progress, 0 failed.
Locality summary:
    distributed : 903 virtual-volumes.
    local      : 1074 virtual-volumes.
    remote    : 175 virtual-volumes.
Cache-mode summary:
    asynchronous : 0 virtual-volumes.
    synchronous  : 2152 virtual-volumes.
Total virtual-volume capacity is 43.7T.
    
```



```

Virtual-volume health summary (cluster-2):
    Total 1991 virtual-volumes, 0 unhealthy.
Expansion summary:
virtual-volume name  expandable-capacity  capacity  expansion-status
-----
dr_one2one_CX_0_vol  10G                10G      -
    Total 1 expansion: 0 in-progress, 0 failed.
Locality summary:
    distributed : 903 virtual-volumes.
    local      : 960 virtual-volumes.
    remote     : 128 virtual-volumes.
Cache-mode summary:
    asynchronous : 0 virtual-volumes.
    synchronous  : 1991 virtual-volumes.
Total virtual-volume capacity is 43.9T.

```

See also

- ds summary
- export port summary
- export storage-view summary
- extent summary
- local-device summary

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
- vpn start
- vpn stop

vpn start

Starts the VPN connection between the management servers.

Contexts

All contexts.

Syntax

```
vpn start
```

Description

Starts the VPN connection between the management servers. In a VPLEX Metro system, run this command on both the clusters.

Example

```
VPlexcli:/> vpn start
```

See also

- `vpn restart`
- `vpn status`
- `vpn stop`

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.

Examples

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
```

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

- `vpn restart`
- `vpn start`
- `vpn stop`

vpn stop

Stops the VPN connection between management servers.

Contexts

All contexts.

Syntax

```
vpn stop
```

Description

Stops the VPN connection between management servers.

Examples

```
Vplexcli: /> vpn stop
```

See also

- `vpn restart`
- `vpn status`
- `vpn start`

wait

Causes a wait until specified context-tree conditions are met.

Contexts

All contexts.

Syntax

```
wait [-c | --context-list] [, context-list ...]
      [-a | --attribute= attribute]
      [-v | --value= value]
      [-t | --timeout= timeout]
      [-h | --help]
      [--verbose]
```

Arguments

Required arguments	
<code>[-c --context-list] [, context-list ...]</code>	Context list, separated by commas

Optional arguments	
[-a --attribute]	Attribute name
[-v --value]	Attribute value
[-t --timeout]	Timeout in seconds. Default is twenty seconds (20s)
[-h --help]	Displays the usage for this command
[--verbose]	Provides additional output during command execution. This may not have any effect for some commands.

Description

If a `context-list` is provided without an attribute, the command will wait until the contexts in the list exist. If wildcard patterns are used, the command will wait until at least one context can be resolved for every pattern.

If an attribute and value pair are given, the command will wait until the attribute of every context resolved from `context-list` has the given value.

The attribute values are compared as strings.

Use the `--timeout` option to set the timeout in seconds. The default timeout is 20 seconds.

webserver

Start, stop, or restart the Webserver.

Contexts

All contexts.

Syntax

```
webserver [stop | start | restart]
          [-h | --help]
          [--verbose]
```

Arguments

Optional arguments	
[-h --help]	Displays the usage for this command
[--verbose]	Provides additional output during command execution. This may not have any effect for some commands.

Description

This command starts, stops, or restarts the Webserver.

Note: To ensure a successful restart of the Webserver, it is recommended to avoid using the `restart` option as it has proven to be unreliable in some cases due to a number of external environmental factors. Instead, to restart, issue a `stop`, and then a `start`. After issuing these commands, verify that the Webserver is running.

INDEX

A

- alias 45
- amp:register 47
- amp:unregister 49
- array fail 51
- array:claim 50
- array:re-discover 52
- array:used-by 54
- audience 13
- authentication directory-
 - service:configure;configure:authentication directory-service 55
- authentication directory-service:map;map:authentication directory-service 59
- authentication directory-service:show;show:authentication directory-service 60
- authentication directory-service:unconfigure 62
- authentication;directory-service unmap 62

B

- Batch migration:perform a migration 71
- batch-migrate check-plan 67
- batch-migrate clean 69
- batch-migrate create-plan 71
- batch-migrate:check-plan;check-plan, batch-migrate 67
- batch-migrate:pause;pause:batch-migrate 73
- batch-migrate:remove;remove:batch-migrate 74
- batch-migrate:resume;resume:batch-migrate 75
- batch-migrate:start;start:batch-migrate 76
- batch-migrate:summary;summary:batch-migrate 78
- battery-conditioning:disable:disable:battery-conditioning 81
- battery-conditioning:enable;enable:battery-conditioning 83
- battery-conditioning:manual-cycle cancel-request 85
- battery-conditioning:manual-cycle request 86
- battery-conditioning:set-enabled 83
- battery-conditioning:set-schedule 87
- battery-conditioning:summary;summary:battery-conditioning 90

C

- cache invalidate 92
- cache-invalidate-status 95
- call-home import-event-modifications 395
- call-home remove-event-modifications 395
- call-home-view-event-modifications 396
- capture:begin;begin capture 98
- capture:end;end capture 99
- capture:pause;pause capture 100
- capture:replay;replay capture 100
- capture:resume;resume capture 101
- cd 102
- chart create 103

- CLI : executing commands;executing CLI commands 28
- CLI :viewing command history; command history, viewing 34
- CLI commands:cd 27
- CLI commands:globbing 33
- CLI commands:ls 27
- CLI commands:popd 27
- CLI commands:positional arguments 33
- CLI commands:pushd 27
- CLI commands:tree 27
- CLI commands:verbose and --help arguments 34
- CLI commands:wildcards 31
- CLI context tree 26
- CLI workspace:command history 34
- CLI:-verbose and --help arguments 34
- CLI:? wildcard 32
- CLI:* wildcard 31
- CLI:** wildcard 31
- CLI:accessing help;help command:CLI ;help, accessing in CLI 35
- CLI:context tree 24, 26
- CLI:display command arguments} 30
- CLI:login 22, 24
- CLI:page output 29
- CLI:positional arguments 33
- CLI:Tab completion 29
- CLI:using commands 28
- CLI:Where am I in the context tree? 27
- cluster show-remote-devices 112
- cluster summary:cluster expelled 119, 123
- cluster-witness:configure 124
- cluster-witness:disable 126
- cluster-witness:enable 128
- cluster:add 105
- cluster:cacheflush 106
- cluster:configdump 107
- cluster:expel 109
- cluster:forget 110
- cluster:shutdown 114
- cluster:status 116
- cluster:summary;summary:cluster 119
- cluster:unexpel 123
- collect diagnostics 133
- commands 17
- comments 13
- configuration sync-time-show 163
- configuration;event-notices-reports config 144
- configuration:complete-system-setup 135
- configuration:configure-auth-service 136
- configuration:connect-local-directors 137
- configuration:connect-remote-directors 138
- configuration:continue-system-setup 139
- configuration:cw-change-password 141
- configuration:cw-uninstall 139
- configuration:cw-vpn-configure 139

- configuration:cw-vpn-reset 141
 - configuration:enable-front-end-ports 143
 - configuration:event-notices-reports-reset 144
 - configuration:event-notices-reports-show 145
 - configuration:flashdir-backup-enable 147
 - configuration:get-product-type 147
 - configuration:join-clusters 148
 - configuration:metadata-backup 149
 - configuration:register-product 151
 - configuration:remote-clusters add-addresses 152
 - configuration:remote-clusters clear-addresses 154
 - configuration:show-meta-volume-candidates 155
 - configuration:subnet clear 156
 - configuration:subnet remote-subnet add 158
 - configuration:subnet remote-subnet remove 159
 - configuration:sync-time 161
 - configuration:sync-time-clear 162
 - configuration:upgrade-meta-slot-count 166
 - configure-certificates 450
 - connect 167
 - connectivity director 169
 - connectivity:list all 170
 - connectivity:list directors 171
 - connectivity:list initiators 172
 - connectivity:list storage-volumes 173
 - connectivity:show;show:connectivity 174
 - connectivity:validate-be;validate-be, connectivity 174
 - connectivity:validate-local-com;local-com:validate connectivity 177
 - connectivity:validate-wan-com;wan-com:validate connectivity 178
 - consistency-group 183
 - consistency-group:add-virtual-volumes 179
 - consistency-group:choose-winner 180
 - consistency-group:create 184
 - consistency-group:destroy 186
 - consistency-group:list-eligible-virtual-volumes 187
 - consistency-group:remove-virtual-volumes 188
 - consistency-group:resolve-conflicting-detach 190
 - consistency-group:resume-at-loser 192
 - consistency-group:set-detach-rule no-automatic-winner 194
 - consistency-group:set-detach-rule winner 195
 - consistency-group:summary 197
 - conventions for publication 13
- D**
- deprecated commands 19
 - describe 198
 - Detach:mirror leg 202
 - device attach-mirror 199
 - device collapse 201
 - device:detach-mirror;detach-mirror, device 202
 - device:mirror-isolation auto-unisolation enable 207
 - device:mirror-isolation auto-unisolation-disable 206
 - device:mirror-isolation disable 209
 - device:mirror-isolation disable;disable:mirror-isolation device 209
 - device:mirror-isolation show 213
 - device:resume-link-down;resume-link-down, device 215
 - device:resume-link-up;resume-link-up, device 217
 - director appcon 219
 - director appdump 220
 - director appstatus 222
 - director commission 223
 - director decommission 224
 - director fc-port-stats 224
 - director firmware show-banks 226
 - director forget 227
 - director shutdown 230
 - director tracepath 231
 - director traceroute 231
 - director:appdump;appdump, director 220
 - director:commission;commission, director 223
 - director:passwd;passwd, director 228
 - director:ping;ping 229
 - director:shutdown;shutdown:director 230
 - dirs 233
 - disconnect 233
 - Display:multiple objects using globs 361
 - Display:object's attributes 361
 - Display:object's attributes and contexts 361
 - Display:operational status 361
 - distributed device 244
 - dm migration pause 237
 - dm:migration cancel;migration:cancel;cancel:dm migration 234
 - dm:migration clean;migration:clean;clean, dm migration 235
 - dm:migration commit;migration:commit;commit, dm migration 236
 - dm:migration remove;migration:remove;remove:dm migration 238
 - dm:migration resume;migration:resume;resume:dm migration 239
 - dm:migration start;migration:start;start:dm migration 240
 - drill-down 242
 - ds dd remove-all-rules 249
 - ds summary 258
 - ds:dd create;distributed-storage:dd create;dr:create;distributed-device:create;create:distributed-device 245
 - ds:dd declare-winner;distributed-device:declare-winner;dr:declare-winner;distributed-storage:dd declare-winner;declare-winner, distributed-device 247
 - ds:dd destroy;distributed-storage:dd destroy;distributed-device:destroy;dr:destroy;destroy:distributed-device 248
 - ds:dd set-log;distributed-storage:dd set-log;distributed-device:set-log;dr:set-log;set-log, distributed-device 250
 - ds:rule destroy;distributed-storage:rule destroy;rule:destroy;destroy:rule 252
 - ds:rule island-containing;distributed-storage:rule island-containing;rule:island-containing;island-containing rule 252
 - ds:rule-set copy;distributed-storage:rule-set copy;rule-set:copy;copy rule-set 254
 - ds:rule-set destroy;distributed-storage:rule-set destroy;rule-set:destroy;destroy:rule-set 255

E

esrs import-certificate 265
 esrs register 266
 esrs status 267
 esrs un-register 268
 event-test 269
 exec 270
 exit 270
 export initiator-port discovery;initiator
 ports:discovering;discovering:initiator ports 271
 export initiator-port:show-logins 274
 export target-port renamewwns 293
 export:initiator-port register-host;register-host, export
 initiator-port;initiator-port:register-host 273
 export:initiator-port register;initiator-
 port:register;register:initiator-port 272
 export:initiator-port unregister;initiator-
 port:unregister;unregister, initiator-port 276
 export:port summary;summary:export port;port, export
 summary 277
 export:target-port renamewwns;target-port,
 renamewwns;renamewwns, export target-port
 293
 export:view
 addinitiatorport;view:addinitiatorport;addinitiator
 port, export view 279
 export:view addport;view:addport;addport, export view
 280
 export:view
 addvirtualvolume;view:addvirtualvolume;addvirtu
 alvolume, export view 280
 export:view checkconfig;checkconfig, export
 view;view:checkconfig 283
 export:view create;view:create;create:export view 284
 export:view destroy;view:destroy;destroy:export view 285
 export:view find-unmapped-volumes;view:find-unmapped-
 volumes;find-unmapped-volumes, export view
 287
 export:view find;find, export view;view:find 286
 export:view map;map, export view;view:map 288
 export:view
 removeinitiatorport;view:removeinitiatorport;rem
 oveinitiatorport, export view 289
 export:view removeport;view:removeport;removeport,
 export view 290
 export:view
 removevirtualvolume;view:removevirtualvolume;r
 emovevirtualvolume, export view 291
 export:view show-powerpath-interfaces;view:show-
 powerpath-interfaces;show-powerpath-
 interfaces, export view 292
 export:view summary;summary:export view;view:summary
 292
 Extent migrations 240
 extent:create;create extent 294
 extent:destroy;destroy:extent 296
 extent:summary;summary:extent 297
 extent:unclaim;unclaim extent 302

F

find 25, 299

front-end-performance-stats start 300
 front-end-performance-stats status 300
 front-end-performance-stats stop 301

G

gateway IP address:set 364
 getsysinfo 302

H

health-check 303
 help 309
 history 310

I

import-event-modifications 395
 IP address:set 364
 IP address:set IP from peer MMCS 364
 iscsi chap back-end:add credentials;iscsi chap:back-end
 add-credentials 311
 iscsi chap back-end:disable;iscsi chap:back-end disable 312
 iscsi chap back-end:enable;iscsi chap:back-end enable 313
 iscsi chap back-end:list-credentials;iscsi chap:back-end
 list-credentials 314
 iscsi chap back-end:remove-credentials 315
 iscsi chap back-end:remove-default credential;iscsi
 chap:back-end remove-default-credential 316
 iscsi chap back-end:set-default-credential;iscsi chap:back-
 end:set-default-credential 317
 iscsi chap front-end:add-credential;iscsi chap:front-end
 add-credential 318
 iscsi chap front-end:disable;iscsi chap:front-end disable
 320
 iscsi chap front-end:enable;iscsi chap:front-end enable 321
 iscsi chap front-end:list-credentials;iscsi chap:front-end
 list-credentials 322
 iscsi chap front-end:remove-credentials;iscsi chap:front-
 end remove-credentials 323
 iscsi chap front-end:remove-default-credential;iscsi
 chap:front-end remove-default-credential 324
 iscsi chap front-end:set-default-credential;iscsi
 chap:front-end set-default-credential 325
 iscsi check-febe-connectivity 326
 iscsi isns add 327
 iscsi isns:list;list:iscsi isns;iscsi:isns list 328
 iscsi isns:remove;remove:iscsi isns;iscsi:isns remove 330
 iscsi sendtargets add 331
 iscsi sendtargets:list;list:iscsi sendtargets;iscsi:sendtargets
 list 332
 iscsi sendtargets:rediscover;iscsi:sendtargets rediscover
 334
 iscsi sendtargets:remove;iscsi:sendtargets
 remove;remove:iscsi sendtargets 335
 iscsi targets:list;list:iscsi targets;iscsi:targets list 336

L

license install 338, 341
 license reset 339
 license show 340
 list-certificates 463

local-device:destroy;destroy:local-device;device:destroy local 344
 local-device:summary;summary:local-device 345
 log filter create 346
 log source create 349
 log source destroy 350
 log:filter create;create:log filter 346
 log:filter destroy;destroy:log filter 348
 log:filter list;list:log filter 348
 log:source create;create:log source 349
 log:source destroy;destroy:log source 350
 log:source list;list:log source 351
 logging-volume detach-mirror 357
 logging-volume:add-mirror;add-mirror, logging-volume 351
 logging-volume:create;create:logging-volume 352
 logging-volume:destroy;destroy:logging-volume 357, 359
 logical-unit forget 360
 ls;list 361

M

management-server:set-ip;set-ip:management-server 364
 manifest:upgrade;upgrade:manifest 366
 manifest:version;version:manifest 366
 meta-volume:attach-mirror;attach-mirror:meta-volume 367
 meta-volume:create;create:meta-volume 369
 meta-volume:destroy;destroy:meta-volume 372
 meta-volume:move;move, meta-volume 374
 metadata backup:display backup volumes 149
 metadata:backup:initial 149
 metadata:backup:modify schedule 149
 metadata:modify backup schedule 149
 metadata:schedule backup 149
 migrations:device 240
 migrations:extent 240
 migrations:prerequisites 240
 mirror leg:detach 202
 monitor add;monitor:add 376
 monitor:add-console-sink; add-console-sink, monitor; sink:add-console 376
 monitor:add-file-sink;add-file-sink, monitor;sink:add-file 378
 monitor:collect;collect:monitor 379
 monitor:create;create:monitor 380
 monitor:destroy;destroy:monitor 383
 monitor:remove-sink;remove-sink, monitor;sink:remove 385
 monitor:stat-list;stat-list, monitor 385

N

Names 32
 Names:VPLEX components 32
 ndu pre-check 388
 ndu pre-config-upgrade 386
 ndu recover 388
 ndu rolling-upgrade c41-to-sles 388
 ndu rolling-upgrade c41x-to-sles 388
 ndu rolling-upgrade ssd-fw 388
 ndu start 391
 ndu status 394
 NDU:display status 394

NDU:pre-config-upgrade 387
 NDU:pre-ndu validation 386
 NDU:recover from failed ndu 388
 NDU:rolling upgrade SSD firmware 388
 NDU:upgrade unconfigured 387
 net-mask:set 364
 Network:set IP address 364
 notifications call-home import-event-modification 395
 notifications call-home test;test, call-home notifications 397
 notifications call-home:import-event-modifications 395
 notifications call-home:remove-event-modifications 395
 notifications call-home:test 397
 notifications call-home:view-event-modifications 396
 notifications job cancel 398
 notifications job delete 398
 notifications job:cancel 398
 notifications job:delete 398
 notifications snmp-trap create 401
 notifications snmp-trap destroy 402
 notifications:call-home remove-event-modifications 395

P

password-policy set 403
 plugin:addurl 405
 plugin:listurl 406
 plugin:register 407
 popd 407
 preface 13
 provisioning 485
 pushd 408

R

rebuild:set-transfer-size 409
 rebuild:show-transfer-size 410
 rebuild:status 411
 related documentation 13
 remove-modifications 395
 report:aggregate-monitors 412
 report:capacity-arrays 413
 report:capacity-clusters 416
 report:capacity-hosts 417
 report:create-monitors 419
 report:poll-monitors 421
 rm 422
 rp:import-certificate 423
 rp:rpa-cluster add 424
 rp:rpa-cluster remove 431
 rp:summary 432
 rp:validate-configuration 434

S

schedule:add 437
 schedule:modify 438
 schedule:remove 439
 scheduleSYR:add 440
 scheduleSYR:list 441
 scheduleSYR:remove 442
 script, command 442
 search 25

security set-login-banner 469, 470
 security:configure-mmcs-users 443
 security:create-ca-certificate 444
 security:create-certificate-subject 446
 security:create-host-certificate 448
 security:delete-ca-certificate 454
 security:delete-host-certificate 455
 security:export-ca-certificate 456
 security:export-host-certificate 457
 security:import-ca-certificate 458
 security:import-host-certificate 459
 security:ipsec-configure 461
 security:remove-login-banner 464
 security:renew-all-certificates 465
 security:show-cert-subj 470
 sessions, command 471
 set 472
 set topology 477
 show-use-hierarchy 478
 shutdown 111, 119
 sms dump 479
 snmp-agent:configure 481
 snmp-agent:start 482
 snmp-agent:status 483
 snmp-agent:stop 483
 snmp-agent:unconfigure 484
 snmp-trap, create;create, snmp-trap 401
 snmp-trap, destroy;destroy, snmp-trap 402
 source 484
 statistics 384
 storage-tool-compose 486
 storage-volume list-thin-capable 500
 storage-volume:auto-unbanish-interval 489
 storage-volume:claim 490
 storage-volume:claimingwizard 493
 storage-volume:find-array 496
 storage-volume:forget 498
 storage-volume:list-banished 499
 storage-volume:resurrect 501
 storage-volume:summary 503
 storage-volume:unbanish 507
 storage-volume:unclaim 508
 storage-volume:used-by 510
 support information 13
 syrcollect 511

T

thin volumes 218
 tree 512

U

unalias 513
 user:remove 519
 user:add 514
 user:event-server add-user 515
 user:event-server change-password 516
 user:list 517
 user:passwd 518
 user:reset 520

V

validate-system-configuration;configuration,
 validate;system, validate configuration 521
 vault:go 522
 vault:overrideUnvaultQuorum 523
 vault:status 525
 verify fibre-channel-switches 529
 version 529
 view-event-modifications 396
 virtual-volume list-thin 545
 virtual-volume set-thin-enabled 548
 virtual-volume:create 533
 virtual-volume:destroy 540
 virtual-volume:expand 541
 virtual-volume:expand (back end) 541
 virtual-volume:provision 546
 virtual-volume:re-initialize 548
 virtual-volume:summary 549
 volume:summary;summary:volume 549
 vpn:restart 553
 vpn:start 553
 vpn:status 554
 vpn:stop 555

W

Wait 555
 web certificate 471
 Webserver 556
 wildcard 25

