



EMC[®] VNX[®] Series

Version 5.33

Command Line Interface Reference for Block

P/N 300-015-135 REV. 06

Copyright © 2013-2016 EMC Corporation. All rights reserved. Published in the USA.

Published September 2016

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

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

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

CONTENTS

Preface		15
Chapter 1	About CLI Commands	19
	naviseccli.....	20
Chapter 2	VNX CLI for Block	23
	Basic commands.....	24
	Storage setup and management overview.....	24
	Shared storage setup and management overview.....	25
	alpa -get.....	26
	alpa -set.....	27
	arraycommpath.....	28
	arrayconfig -capture.....	29
	arrayconfig -duplicate.....	30
	arrayname.....	32
	backendbus -analyze.....	32
	backendbus -get -connstate.....	35
	backendbus -get -physical.....	36
	backendbus -get -prominfo.....	36
	backendbus -get -sfpstate.....	37
	backendbus -get -speeds.....	38
	backendbus -get -all.....	38
	backendbus -resetspeed.....	41
	baseuuid.....	42
	bind.....	43
	cachecard -list.....	47
	cachecard -initialize.....	48
	cache -sp -info.....	49
	cache -sp -modify.....	50
	chglun.....	50
	chrg.....	52
	clearlog.....	53
	clearstats.....	54
	cmdtime.....	54
	copytodisk.....	55
	createrg.....	56
	emconfiguration.....	58
	environment.....	58
	failovermode.....	61
	faults -list.....	62
	flash -ioport.....	63
	flash -iomodule.....	64
	flashleds.....	65
	getagent.....	66
	getall.....	68
	getarrayuid.....	71
	getconfig.....	72
	getcontrol.....	73

getcrus.....	74
getdisk.....	78
getlog.....	84
getloop.....	85
getlun.....	85
getresume.....	92
getrg.....	95
getsniffer.....	97
getsp.....	100
getsptime.....	101
getspuptime.....	101
getunusedluns.....	102
hotsparepolicy -list.....	103
hotsparepolicy -set.....	103
inserttestevent.....	104
ioportconfig -config -maxvlans.....	104
ioportconfig -list.....	105
ioportconfig -list -maxvlans.....	112
ioportconfig -persist.....	112
luncache -clear.....	113
luncache -list.....	114
managedby.....	114
managefiles -delete.....	115
managefiles -list.....	116
managefiles -retrieve.....	116
ndu -list.....	117
ndu -install.....	119
ndu -runrules.....	122
ndu -commit.....	123
ndu -status.....	124
ndu -revert.....	124
networkadmin -get.....	125
networkadmin -mib.....	129
networkadmin -route.....	130
networkadmin -set.....	131
port.....	134
powersaving.....	145
rebootSP.....	146
rebootpeerSP.....	147
remoteconfig -getconfig.....	147
remoteconfig -setconfig.....	148
removerg.....	150
resetandhold.....	150
responsetest.....	151
sc_off.....	153
setfeature.....	154
setsniffer.....	154
setsniffverify.....	155
setspstime.....	156
setstime.....	156
setstats.....	157
shutdown.....	158
shutdownsp.....	159
shutdownpeersp.....	159
spcollect -info.....	160
spcollect -set.....	160

spcollect.....	161
spportspeed.....	162
storagegroup.....	168
systemtype.....	179
trespass.....	180
unbind.....	181
unitserialnumber.....	181
LUN Expansion (metaLUN) commands.....	182
MetaLUNs overview.....	182
CLI metaLUN command functions.....	183
metalun -destroy.....	184
metalun -expand.....	184
metalun -info.....	190
metalun -list.....	191
metalun -modify.....	196
Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN.....	197
Reserved LUN Pool commands.....	198
reserved -lunpool -addlun.....	199
reserved -lunpool -list.....	200
reserved -lunpool -rmlun.....	201
iSCSI and FCoE commands.....	202
iSCSI commands overview.....	202
connection -adduser.....	203
connection -deleteuser.....	204
connection -getuser.....	205
connection -getport.....	206
connection -setport.....	217
connection -delport.....	219
connection -route.....	220
connection -pingnode.....	221
connection -traceroute.....	222
connection -setsharedauth.....	223
connection -delsharedauth.....	224
connection -getsharedauth.....	224
connection -addset.....	225
connection -delset.....	226
connection -modifyset.....	226
connection -getset.....	228
connection -addpath.....	228
connection -delpath.....	229
connection -modifypath.....	230
connection -verifypath.....	231
isns -addserver.....	232
isns -deleteserver.....	232
isns -listserver.....	233
isns -setprimary.....	234
LUN Migration commands.....	235
migrate -start.....	235
migrate -cancel.....	236
migrate -modify.....	237
migrate -list.....	238
Domain and Security commands.....	239
About domains and user accounts.....	239
Domain and security commands overview.....	242
domain -add.....	243

domain -list.....	245
domain -remove.....	246
domain -uninitialize.....	247
domain -setmaster.....	248
ntp -list.....	248
ntp -set.....	250
security -adduser.....	251
security -changeuserinfo.....	253
security -list.....	254
security -listrole.....	254
security -rmuser.....	255
security -ldap -addserver.....	255
security -ldap -modifyserver.....	257
security -ldap -removeserver.....	259
security -ldap -listserver.....	260
security -ldap -addrolemapping.....	261
security -ldap -modifyrolemapping.....	261
security -ldap -removerolemapping.....	262
security -ldap -listrolemapping.....	263
security -ldap -synchronize.....	263
security -certificate -add.....	264
security -certificate -remove.....	264
security -certificate -list.....	265
security -certificate -alertBeforeExpiration.....	266
security -pkcs12upload.....	266
security -SPcertificate -generate.....	267
security -fipsmode -set.....	267
security -fipsmode -get.....	268
security -setport.....	268
security -trustedclient -enable.....	268
security -trustedclient -add.....	269
security -trustedclient -list.....	270
security -trustedclient -remove.....	270
Event Monitor commands.....	271
Event Monitoring overview.....	271
-portal -create.....	272
-portal -list.....	272
-portal -migrate.....	273
-portal -destroy.....	273
-template -list.....	274
-template -import.....	274
-template -export.....	275
-template -destroy.....	275
-template -findconflict.....	276
-template -swap.....	276
-monitor -addsystem.....	277
-monitor -removesystem.....	277
-monitor -listmapping.....	278
-monitor -applytemplate.....	278
-monitor -stoptemplate.....	279
-monitor -getlog.....	280
-monitor -enableresponse.....	280
-monitor -disableresponse.....	281
-monitor -getlogsize.....	282
-monitor -setlogsize.....	282
-monitor -inserttestevent.....	283

	-monitor -responsetest.....	283
	-monitor -reloadconfig.....	285
Virtual server commands.....		285
Virtual server overview.....		285
server -volmap.....		286
server -update.....		288
server -getagent.....		288
server -remoteconfig.....		289
server -register.....		290
hypervisor -addvirtualcenter.....		290
hypervisor -modifyvirtualcenter.....		291
hypervisor -removevirtualcenter.....		292
hypervisor -addesx.....		292
hypervisor -modifyesx.....		293
hypervisor -removeesx.....		293
hypervisor -listesx.....		294
hypervisor -bypassCert -set.....		294
hypervisor -bypassCert -get.....		294
Host Agent commands.....		295
server -volmap.....		295
server -update.....		297
server -getagent.....		297
server -remoteconfig.....		298
server -register.....		299
Unisphere Server Utility.....		299
Using the server utility command line to generate a high-availability report.....		299
Unisphere Initialization Utility.....		302
Using the command line initialization utility.....		302
Chapter 3	Virtual Pools Commands	305
Virtual Pools overview.....		306
storagepool -create.....		307
storagepool -list.....		310
storagepool -destroy.....		316
storagepool -modify.....		316
storagepool -expand.....		318
storagepool -cancelexpand.....		319
storagepool -feature -info.....		319
storagepool -resumeharvesting.....		321
lun -create.....		321
lun -modify.....		323
lun -expand.....		325
lun -destroy.....		326
lun -list.....		326
deduplication -feature -pause.....		330
deduplication -feature -resume.....		330
deduplication -feature -info.....		330
Chapter 4	Auto-tiering Commands	333
About tiered storage.....		334
autotiering -schedule -modify.....		335
autotiering -schedule -enable -disable.....		336
autotiering -relocation -start.....		336

	autotiering -relocation -setrate.....	337
	autotiering -relocation -pause resume.....	338
	autotiering -relocation -stop.....	338
	autotiering -info.....	339
Chapter 5	Compression Commands	341
	About LUN Compression.....	342
	compression -on.....	342
	compression -pause.....	343
	compression -feature -pause.....	344
	compression -resume.....	344
	compression -feature -resume.....	345
	compression -feature -info.....	345
	compression -modify.....	345
	compression -off.....	346
	compression -list.....	346
Chapter 6	FAST Cache Commands	349
	Introduction to caching.....	350
	cache -fast -create.....	351
	cache -fast -destroy.....	352
	cache -fast -info.....	352
Chapter 7	EMC Unisphere QoS Manager	355
	About Unisphere Quality of Service Manager.....	356
	QoS Manager prerequisites.....	356
	QoS Manager CLI commands.....	356
	nqm -getlatestperf.....	357
	nqm -info.....	357
	nqm -ioclass -create.....	358
	nqm -ioclass -destroy.....	359
	nqm -ioclass -list.....	360
	nqm -ioclass -modify.....	362
	nqm -measure.....	364
	nqm -off.....	364
	nqm -policy -create.....	365
	nqm -policy -destroy.....	366
	nqm -policy -list.....	366
	nqm -policy -modify.....	367
	nqm -run.....	368
	nqm -schedule -create.....	369
	nqm -schedule -destroy.....	370
	nqm -schedule -list.....	370
	nqm -schedule -modify.....	372
	nqm -schedule -resume.....	373
	nqm -schedule -suspend.....	373
	nqm -setfallback.....	374
	nqm -setoptions.....	374
	nqm -stop.....	375
	Client tools.....	375
	nqm -archive -dump.....	375
	nqm -archive -retrieve.....	382

Chapter 8	EMC Unisphere Analyzer	385
	About Unisphere Analyzer.....	386
	Prerequisites.....	386
	Analyzer CLI commands.....	386
	analyzer -set.....	387
	analyzer -get.....	388
	analyzer -logging.....	388
	analyzer -start.....	389
	analyzer -stop.....	389
	analyzer -status.....	390
	analyzer -archive.....	390
	Client tools.....	392
	analyzer -archivedump.....	392
	analyzer -archiveretrieve.....	405
	analyzer -archivemerge.....	406
Chapter 9	EMC VNX MirrorView/A CLI	409
	MirrorView/A remote mirroring commands.....	410
	MirrorView/A operations overview.....	410
	MirrorView/A CLI functions.....	412
	mirror -async -addimage.....	413
	mirror -async -change.....	415
	mirror -async -changeimage.....	416
	mirror -async -create.....	418
	mirror -async -destroy.....	419
	mirror -disablepath.....	420
	mirror -enablepath.....	421
	mirror -async -fractureimage.....	422
	mirror -async -info.....	423
	mirror -async -list.....	425
	mirror -async -promoteimage.....	427
	mirror -async -removeimage.....	428
	mirror -async -setfeature.....	429
	mirror -async -syncimage.....	430
	MirrorView/A consistency group commands.....	431
	mirror -async -creategroup.....	431
	mirror -async -destroygroup.....	433
	mirror -async -addtogroup.....	434
	mirror -async -removefromgroup.....	434
	mirror -async -changegroup.....	435
	mirror -async -syncgroup.....	437
	mirror -async -fracturegroup.....	437
	mirror -async -promotegroup.....	438
	mirror -async -listgroups.....	439
Chapter 10	EMC VNX MirrorView/S CLI	443
	MirrorView/S Remote Mirroring commands.....	444
	MirrorView/S operations overview.....	444
	MirrorView/S CLI functions.....	446
	mirror -sync -addimage.....	447
	mirror -sync -allocatelog.....	450
	mirror -sync -change.....	450
	mirror -sync -changeimage.....	452

mirror -sync -create.....	453
mirror -sync -deallocatelog.....	455
mirror -sync -destroy.....	455
mirror -disablepath.....	457
mirror -enablepath.....	458
mirror -sync -fractureimage.....	459
mirror -sync -info.....	460
mirror -sync -list.....	462
mirror -sync -listsyncprogress.....	464
mirror -sync -listlog.....	465
mirror -sync -promoteimage.....	465
mirror -sync -removeimage.....	467
mirror -sync -setfeature.....	469
mirror -sync -syncimage.....	470
MirrorView/S consistency group commands.....	471
mirror -sync -creategroup.....	471
mirror -sync -destroygroup.....	472
mirror -sync -addtogroup.....	473
mirror -sync -removefromgroup.....	474
mirror -sync -change group.....	475
mirror -sync -syncgroup.....	476
mirror -sync -fracturegroup.....	477
mirror -sync -promotegroup.....	478
mirror -sync -listgroups.....	479

Chapter 11 SAN Copy and its command line utilities 481

Introducing SAN Copy software.....	482
Preparing a Block system for a copy session.....	482
Connecting SAN Copy ports to storage groups.....	483
Preparing a Symmetrix system for a copy session.....	485
Preparing a Symmetrix system using EMC ControlCenter SAN Manager.....	486
Preparing a Symmetrix system using ESN Manager.....	487
About admhost for SAN Copy.....	488
admhost SAN Copy command functions.....	488
Preparing logical units for a SAN Copy session.....	490
Preparing source and destination logical units for full and incremental sessions.....	491
SAN Copy CLI commands.....	492
SAN Copy command functions.....	492
sancopy -create.....	494
sancopy -destinfo.....	498
sancopy -duplicate.....	501
sancopy -info.....	502
sancopy -mark.....	506
sancopy -modify.....	507
sancopy -pause.....	512
sancopy -remove.....	513
sancopy -resume.....	514
sancopy -settings.....	515
sancopy -start.....	516
sancopy -stop.....	518
sancopy -throttle.....	519
sancopy -transfer.....	520
sancopy -unmark.....	521

	sancopy -updateconnections.....	521
	sancopy -verify.....	522
Chapter 12	EMC SnapView Software	525
	SnapView clone CLI commands.....	526
	Setting up SnapView to use clones.....	526
	CLI clone command functions.....	527
	clone snapview -addclone.....	528
	clone snapview -allocatcpl.....	531
	clone snapview -changeclone.....	532
	clone snapview -changeclonefeature.....	534
	clone snapview -changeclonegroup.....	535
	clone snapview -consistentfractureclones.....	536
	clone snapview -createclonegroup.....	537
	clone snapview -deallocatcpl.....	538
	clone snapview -destroyclonegroup.....	539
	clone snapview -fractureclone.....	540
	clone snapview -listclone.....	542
	clone snapview -listclonefeature.....	544
	clone snapview -listclonegroup.....	546
	clone snapview -removeclone.....	548
	clone snapview -resetfracturelog.....	549
	clone snapview -reversesyncclone.....	550
	clone snapview -syncclone.....	552
	Admsnap commands for clones.....	554
	About admsnap software and clones.....	554
	Requirements for admsnap and clones.....	554
	admsnap clone_activate.....	555
	admsnap clone_deactivate.....	556
	admsnap flush.....	558
	admsnap list.....	559
	SnapView Snapshot CLI commands.....	560
	Setting up SnapView to use snapshots.....	560
	SnapView command functions.....	562
	snapview -activatesnapshot.....	563
	snapview -chgrollback.....	565
	snapview -chgsnapshot.....	566
	snapview -createsnapshot.....	566
	snapview -deactivatesnapshot.....	568
	snapview -listrollback.....	569
	snapview -listsessions.....	571
	snapview -listsnapshots.....	573
	snapview -listsnapableluns.....	574
	snapview -lunpool.....	575
	snapview -rmsnapshot.....	576
	snapview -startrollback.....	577
	snapview -startsession.....	579
	snapview -stopsession.....	582
	Admsnap commands for snapshots.....	583
	About admsnap software and snapshots.....	583
	Requirements for snapshots.....	584
	admsnap activate.....	584
	admsnap deactivate.....	588
	admsnap flush.....	591
	admsnap list.....	593

	admsnap start.....	594
	admsnap stop.....	598
Chapter 13	EMC VNX Snapshots Software	601
	VNX Snapshots CLI commands.....	602
	snap -feature.....	602
	snap -create.....	602
	snap -destroy.....	604
	snap -list.....	604
	snap -modify	606
	snap -copy.....	607
	snap -restore.....	608
	snap -attach.....	609
	snap -detach.....	609
	snap -group -create.....	609
	snap -group -destroy	610
	snap -group -modify.....	610
	snap -group -list.....	611
	snap -group -addmember.....	612
	snap -group -rmmember.....	612
	snap -group -replmember.....	613
	lun -create.....	613
	lun -create.....	614
	lun -destroy.....	614
	lun -modify.....	615
	lun -expand.....	616
	lun -attach.....	616
	lun -detach.....	617
	storagepool -create.....	617
	storagepool -modify.....	618
	storagepool -list.....	619
	storagepool -resumeAutoDelete.....	621
	getlun.....	621
	SnapCLI commands for VNX Snapshots.....	623
	snapcli create	623
	snapcli destroy.....	624
	snapcli copy.....	625
	snapcli attach.....	626
	snapcli detach.....	627
Chapter 14	EMC Secure Remote Support on the Storage Processor	629
	ESRS Device Client on Storage Processor.....	630
	esrsconfig -agentprovision	630
	esrsconfig -agentstatus.....	631
	esrsconfig -configproxy -addproxy.....	632
	esrsconfig -configproxy -list.....	632
	esrsconfig -configproxy -removeproxy.....	633
	esrsconfig -configpolicymanager -addpm.....	633
	esrsconfig -configpolicymanager -list.....	634
	esrsconfig -configpolicymanager -addpolicyproxy.....	634
	esrsconfig -configpolicymanager -removepolicyproxy.....	635
	esrsconfig -configpolicymanager -removePM.....	635

Chapter 15	Data at Rest Encryption	637
	Data at Rest Encryption overview.....	638
	securedata -feature -activate.....	639
	securedata -feature -info.....	640
	securedata -backupkeys -status.....	641
	securedata -backupkeys -retrieve.....	642
	securedata -auditlog -retrieve.....	643
	securedata -auditlog -cksum.....	644

CONTENTS

Preface

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

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

Audience

This guide is part of the EMC VNX documentation set, and is intended for use by those who will use CLI commands to configure, manage, or obtain the status of EMC systems.

This version of the manual covers up to VNX systems running release 1.3 of Unisphere.

Readers of this guide are expected to be familiar with the following topics:

- The operating system running on the servers you will manage.
- System components and configurations.

For introductory information, refer to the appropriate configuration planning guide for your system.

Special notice conventions used in this document

EMC uses the following conventions for special notices:



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



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



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



Addresses practices not related to personal injury.

Note

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

Systems no longer covered in this document

The table below lists the systems that are no longer covered in this document and the last revision of this document that included the systems.

System removed	Last revision including the system
CX200, CX400, and CX600	300-003-628 Rev. A07 069001184 Rev. A11 300-001-335 Rev. A11 300-004-210 Rev. A02 069001181 Rev. A14
FC series; C series	069001038 Rev. A14 300-001-335 Rev. A11 069001184 Rev. A11
FC4700	069001184 Rev. A08

Related documentation

Related documentation for EMC Navisphere CLI reference includes:

We recommend that you download the latest information before you run the CLI commands. Platform-specific limitations and requirements are provided in the appropriate release note documents.

- *EMC Unisphere Release Notes*
- *EMC MirrorView/Asynchronous and MirrorView/Synchronous for VNX OE for Block Release Notes*
- The version of this manual that is applicable to your software revision

The most up-to-date information about the EMC SAN Copy software is posted on EMC Online Support. We recommend that you download the latest information before you start the SAN Copy software. If you purchased this product from an EMC reseller and you cannot access EMC Online Support, the latest product information should be available from your reseller.

To access EMC Online Support, use the following link:

<http://Support.EMC.com>

Where to get help

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

Product information—For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to EMC Online Support (registration required) at <http://Support.EMC.com>.

Troubleshooting—Go to EMC Online Support at <http://Support.EMC.com>. After logging in, locate the applicable Support by Product page.

Technical support—For technical support and service requests, go to EMC Customer Service on EMC Online Support at <http://Support.EMC.com>. After logging in, locate the applicable Support by Product page, and choose either **Live Chat** or **Create a service request**. To open a service request through EMC Online Support, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Note

Do not request a specific support representative unless one has already been assigned to your particular system problem.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications.

Please send your opinion of this document to:

`techpubcomments@EMC.com`

CHAPTER 1

About CLI Commands

This chapter explains the EMC VNX CLI for Block commands.

Major topic include:

- [navisecli](#)..... 20

naviseccli

Sends status or configuration requests to a system through the command line.

PREREQUISITES

Anyone that can log in to the server running VNX CLI for Block 4.30 or later.

DESCRIPTION

The `naviseccli` command sends system management and configuration requests to a system through the Internet.

SYNTAX

```
naviseccli -help
```

```
naviseccli [-address IPAddress | NetworkName | -h IPAddress |
NetworkName]
[-AddUserSecurity]
[-f filename]
[-nopoll | -np]
[-parse | -p]
[-password password]
[-port port]
[-q]
[-RemoveUserSecurity]
[-scope 0 | 1 | 2]
[-secfilepath]
[security -certificate]
[-timeout | -t timeout]
[-user username]
[-v]
[-xml]
CMD <optional command switches>
```

OPTIONS

`-help`

Displays the help screen and does not start the `naviseccli` process. To start the `naviseccli` process, use one or more of the switches that follows instead.

`-address IPAddress | NetworkName | -h IPAddress | NetworkName`

Specifies the IP address or network name of the targeted SP on the desired system. The default, if you omit this switch, is `localhost`.

`-AddUserSecurity`

Directs the CLI to add user security information to the security file on this server. You must use the `-scope` switch to add scope information to the security file. You can use the `-password` switch or enter your password into the password prompt, to supply the required password information to the security file. The `-user` and `-secfilepath` switches are optional with this command.

Note

If you specify the `-user` switch, you can create an alternative username to your server login name in the security file you create on this server. If you use the `-secfilepath` switch, you can specify an alternative location to your default home directory, for the security file on this server. You must then use the `-secfilepath` switch in each subsequent command that you issue.

`-f filename`

Stores the data in a file.

`-nopoll | -np`

Directs the feature provider not to issue a poll request. This switch significantly increases performance when dealing with large or multiple systems. The feature provider automatically issues a poll request unless this switch is specified.

Note

When the `-nopoll` switch is set, `get` commands may return stale data and `set` commands may erase previously changed settings. Use caution when the `-nopoll` switch is set.

`-parse | -p`

Directs the CLI to validate the command. The CLI verifies the command syntax and displays a message stating whether the command was valid. The CLI takes no other action.

`-password password`

Specifies the password on the system you want to log in to. The password is visible in the command line. Passwords are case-sensitive. If you want to mask the password, and you are not using a security file, you can omit this switch from the command line. The CLI then prompts you to enter a password. The information you enter into the password prompt is concealed.

Note

You can omit this switch if you are using a security file. See `-AddUserSecurity`.

`-port portnumber`

Sets the port number (type) of the system. The default is 443. If you choose to change the default port number, management port 2163 will be supported; however, you will need to specify the `-port` switch and number 2163 in every subsequent command you issue.

`-q`

Suppresses error messages. This switch is useful when included as part of a script.

`-RemoveUserSecurity`

Directs the CLI to remove user security information about the current user from the security file on this server.

`-scope 0 | 1 | 2`

Specifies whether the user account on the system you want to log in to is local, global, or lightweight directory access protocol (LDAP). A 0 (default) indicates global, 1 indicates local, and 2 indicates LDAP. A global account is effective throughout the domain. When the administrator creates a global account, the software copies the definition of this account to the domain directory, which makes it accessible on all systems in the domain. A local account is effective only on the systems for which the administrator creates the account. You can log in only to those systems on which you have a local account. LDAP maps the username or password entries to an

external LDAP or active directory server for authentication. Username or password pairs whose roles are not mapped to the external directory will be denied access.

`-secfilepath filename`

Stores the security file in a file path location you specify. When you create a security file on a server using the `-addusersecurity` command, the security file is saved to your default home directory. If you want to store the security file in an alternative location, you can use the optional `-secfilepath` switch with the `-addusersecurity` command.

Note

If you use the `-secfilepath` switch to set up an alternative path for your security file, you must use this switch in every subsequent CLI command you issue, to ensure the CLI locates the security file.

`security -certificate`

Saves the certificates to the certificate store. It also provides options like `-list`, `-remove`, `-add` to manage the certificate store.

`-timeout | -t timeout`

Sets the timeout value in seconds. The default is 600 seconds.

`-user username`

Specifies the username on the system you want to log in to. Username is case-sensitive. You can omit this if you have added the username to the security file.

Note

You can use this switch when establishing a security file, to specify an alternative username. See `-AddUserSecurity`.

`-v`

Enables verbose error descriptions. This is the default unless `-q` is specified.

`-xml`

Specifies command output in XML format. Use the `-o` (override switch) when specifying `-xml` with commands that require confirmation. Otherwise, the XML output will contain your confirmation string.

CMD

One of the set of commands used with the `naviseccli` command.

CHAPTER 2

VNX CLI for Block

This chapter includes content for VNX series systems.

Major topics include:

- [Basic commands](#)..... 24
- [LUN Expansion \(metaLUN\) commands](#)..... 182
- [Reserved LUN Pool commands](#)..... 198
- [iSCSI and FCoE commands](#)..... 202
- [LUN Migration commands](#)..... 235
- [Domain and Security commands](#)..... 239
- [Event Monitor commands](#)..... 271
- [Virtual server commands](#)..... 285
- [Host Agent commands](#)..... 295
- [Unisphere Server Utility](#)..... 299
- [Unisphere Initialization Utility](#)..... 302

Basic commands

This section explains each of the basic variations, that is, the CLI commands that are common to all systems. Command variations for Storage Groups (Access Logix) are explained in the following chapter.

Storage setup and management overview

This section shows how to configure and monitor a system using the `naviseccli` command. It shows commands with UNIX device names. The Windows device names would be drive letters. In this system, in which each SP is a host. For an overview showing systems with storage groups, see the section [Shared storage setup and management overview on page 25](#). For systems without storage groups:

Procedure

1. Create three RAID groups.

```
naviseccli -h ssl_spa createrg 0 0_0_4 0_0_5 0_0_6 0_0_7 0_0_8 -
raidtype r5
```

Creates RAID group 0 from disks 4 through 8 in the DPE enclosure.

```
naviseccli -h ssl_spa createrg 1 0_0_9 0_0_10 -raidtype r1
```

Creates RAID group 1 from disks 9 and 10 in the DPE enclosure.

```
naviseccli -h ssl_spa createrg 2 0_0_11 0_0_12 -raidtype r1
```

Creates RAID group 2 from disks 11 and 12 in the DPE enclosure.

2. Bind a LUN on each RAID group.

```
naviseccli -h ssl_spa bind 0 -rg 0
```

This command binds a LUN of type RAID 5 with LUN ID 0 on RAID group 0. The LUN occupies all space on RAID group 0, since the bind command did not include the `-cap` switch. By default, read and write caching are enabled on the new LUN.

```
naviseccli -h ssl_spa bind 1 -rg 1
```

This command binds a LUN of type RAID 1 (mirrored pair) with LUN ID 1 on RAID group 1. The LUN occupies all space on RAID group 1 and caching is enabled by default.

```
naviseccli -h ssl_spa bind 2 -rg 2
```

This command binds a LUN of type RAID 1 (mirrored pair) with LUN ID 2 on RAID group 2. The LUN occupies all space on RAID group 2, and caching is enabled by default.

3. Get SP memory information to prepare to set up the system cache.

```
naviseccli -h ssl_spa getsp -mem
SP A
Memory Size For The SP: 12288
SP B
Memory Size For The SP: 12288
```


4. For easy identification in the future, name the system.

```
naviseccli -h ssl_spa arrayname users11
Change the arrayname from xxx to users11? (y/n) y
```

This command changes the name of the system to users11.

Results

At this point, you can create file systems on the LUNs from the operating system and store data on them.

You can learn LUN device names using the `getagent` command and check status periodically using the `getsp`, `getdisk`, and `getlun` commands.

Shared storage setup and management overview

Storage groups provide a simple means of managing which hosts can access defined groups of LUNs. Specify host-storage group access using the unique ID (also called the world wide name, WWN) of each HBA and LUN.

Generally, you may find it easier to use Unisphere than the CLI to create and manipulate storage groups. However, the following is a sequence of tasks you can perform using only the CLI:

Procedure

1. Plan the LUNs that each host will use: RAID group type, capacity, SP owner, storage group, and so on.
2. Create the RAID groups, bind the LUNs, and assign the LUNs to RAID groups using manager or the CLI. See the unshared storage overview in [Storage setup and management overview on page 24](#) for example use of the CLI.

Note

If you are running VMware ESX Server, you can bind LUNs from ESX Server or a virtual machine (VM).

3. Create the storage groups you want using Unisphere or the CLI command `storagegroup -create`.
4. Connect the hosts to storage groups using Unisphere or the CLI. Use the CLI command `storagegroup -connecthost` to assign a storage group to all HBAs in a host, or the command `storagegroup -setpath` to assign a storage group to a specific HBA in a host.

Note

If you are running VMware ESX Server, assign ESX Server to the storage group.

5. Assign LUNs to the storage groups using Unisphere or the CLI `storagegroup -addhlu` command. You can assign more than one LUN to a storage group and let more than one host access a storage group.

Note

If you are running VMware ESX Server, rescan the bus at the ESX Server level. If you are running ESX Server 2.5.x, you must then power down any virtual machine to which you will assign LUNs. If you are running ESX Server 3.x, you do not need to power down the virtual machine. At the ESX Server level, assign the LUNs to the virtual machines. When you have finished, if you are running ESX Server 2.5.x, power up the VM. If you are running ESX Server 3.x, you do not need to power up the VM.

- As needed, reconfigure host-storage group connections using Unisphere or the CLI command `storagegroup -sethost`. If you need to disconnect a host from a storage group, use `storagegroup -disconnecthost`.

alpa -get

Displays the SCSI ID associated with an SP port.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The CLI `alpa` commands (Arbitrated Loop Physical Address) `get` and `set` the port ID SCSI IDs on an SP. These IDs are required for I/O to the SP.

We suggest you use a unique SCSI ID for each SP port in your installation. For example, on the first system, for ports 0 and 1, you can specify SCSI IDs 0 and 1, respectively. On the second system, for the ports you can specify IDs 2 and 3 respectively, and so on.

The `naviseccli` command `-get` function lists ALPA information. The information includes the port ID, the SP ID (A or B), and the SCSI ID associated with the port. Issue the command to the SP for which this information is needed.

If you really need to change an SP port ID, you can do so with the `alpa -set` function.

SYNTAX

```
alpa -get
```

EXAMPLE

For SP A (SP hostname `payroll_storage_spa`), this command lists the SCSI ID associated with port 0.

```
naviseccli -h payroll_storage_spa alpa -get
```

```
Storage Processor: SP A
Port ID:          0
ALPA Value:      0
```

OUTPUT

See above. If the version of VNX Operating Environment (OE) for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
VALID_VALUES_0_3
```

alpa -set

Updates the SCSI ID associated with an SP port.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli alpa` command with the `-set` switch updates the SCSI ID value for the given port on an SP.

⚠ CAUTION

The SP port ID SCSI IDs are initially set by EMC service personnel to work at your site. Do not change any value unless you are installing a new SP and need to change its SCSI IDs from the SP ship values of 0 and 1.

If you change any value, after you confirm, the SP will restart and use the new values.

The software will not let you select a SCSI ID out of range (0-126) or a duplicate ID on a system. If you omit the `-o` (override) switch, then the CLI prompts for confirmation:

```
Changing the ALPA value of port port-id from old-SCSI-ID to new-SCSI-ID (y/n)
```

To confirm, answer **y**; to take no action, answer **n**.

SYNTAX

```
alpa -set
-sp a|b -portid portid SCSI-ID [-o]
```

OPTIONS

- `-sp a|b`
Specifies the SP: A or B.
- `-portid portid SCSI-ID`
Specifies the new SCSI ID to associate with the port. Valid port SCSI IDs range from 0 through 126.
- `-o`
Executes the command without prompting for confirmation.

EXAMPLE

For SP A, this command changes the SCSI ID associated with port 1 to 1.

```
naviseccli -h ss1_SPA alpa -set -sp a -portid 1 1
```

```
Changing the ALPA value of port 1, SP A to 1
WARNING: This operation will reboot the target SP.
Verify that the SPS on the surviving SP is operating correctly.
Otherwise write cache must be disabled:
naviseccli -h IPAddress cache -sp -modify -wc off.
Failure to disable write cache may result in the loss of data.
DO YOU WISH TO CONTINUE? (y/n)
```

OUTPUT

See above. If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
VALID_VALUES_0_126
VALID_VALUES_0_3
```

arraycommpath

Sets a communication path to a system that has no LUNs bound, or removes such a path (non-storage group only).

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli arraycommpath` command creates or removes a communication path between the server and system. Generally this command is needed when you want to configure a system that has no LUNs bound.

This command works only for systems without storage groups. For systems with storage groups, use the command `storagegroup sethost -arraycommpath` (refer [storagegroup on page 168](#).)

The CLI will not issue a confirmation request first if you include the `-o` switch.

⚠ CAUTION

Do not issue this command unless you know the appropriate value for the type of failover software running on the attached servers. The command is designed for transition from one type of failover software to another. Changing the `arraycommpath` mode to the wrong value will make the system inaccessible.

To discover the current setting of `arraycommpath`, use the command without an argument.

SYNTAX

```
arraycommpath [0 | 1] [-o]
```

OPTIONS

- 0 Disables a communication path.
- 1 Enables a communication path.
- `-o` Executes the command without prompting for confirmation.

EXAMPLE

This command disables the communication path between server and system. In this example, the confirmation message shows that the change, if confirmed, will cause the previous setting to be lost.

```
naviseccli -h ssl_spa arraycommpath 0
```

```
WARNING: Previous arraycommpath setting will be lost!
DO YOU WISH TO CONTINUE? (y/n) y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

Valid values are 0 or 1

arrayconfig -capture

Queries the system for configuration information.

PREREQUISITES

A user role of administrator is required to execute this command.

DESCRIPTION

The `arrayconfig -capture` command queries the system for its configuration along with I/O port configuration information. When issued, the command will capture a system's essential configuration data. The information is formatted and stored on the client workstation.

This generated file can be used as a template to configure other systems or rebuild the same system if the previous configuration is destroyed. The information captured may include:

- Software installed in the system (name of the software, revision, and so on).
- Physical configuration of the system (storage processors, disks, enclosures, and so on.)
- Logical configuration of the system (RAID groups, LUNs, storage groups, clones, SnapView, and mirrors).
- I/O port configuration information (hardware and logical configuration information) of a VNX series system.

SYNTAX

```
arrayconfig -capture [-format XML|CSV] [output path to local file name] [-logprogress] [-o]
```

OPTIONS

`-format XML|CSV`

Defaults to XML if `-format` is not specified.

`-logfile`

Retrieves and outputs the log file stored on the targeted system.

`-o`

Executes the command without prompting for confirmation.

`-output path to local file name`

An XML file in the specified path and name will be created. If the indicated output file already exists, you will be asked if you wish to overwrite it.

`-logprogress`

Displays the detailed information. If this option is not specified, displays only error information.

EXAMPLE

```
naviseccli -h ss1_spa arrayconfig -capture -output c:\Capture158.xml
-format xml
```

OUTPUT

If the command succeeds, the XML file will be generated. There will be no message displayed in the command line.

arrayconfig -duplicate

Creates the components in the indicated configuration file on the targeted system.

PREREQUISITES

Manager privileges are required to execute the `arrayconfig -duplicate` command.

DESCRIPTION

Issuing the `arrayconfig -duplicate` command creates the logical components listed in the configuration file (generated from issuing the `arrayconfig -capture` command) for the indicated system. You can view both the status of the duplication operation and the log file.

Before starting the duplication process, the following will be checked to ensure the targeted system qualifies for duplication.

Physical components:

- The model number on the targeted system must be exactly the same as the source.
- The targeted system must have at least the same number of disks as the source (includes all RAID groups).
- The disk enclosure numbers and slot numbers should exactly match the source (includes all RAID groups).
- Drive types must match.
- The size of the targeted disks must be equal to or greater than the size of the source disks.

Software availability:

- The targeted system must have all the system and Unisphere software packages available in the source system.

Logical components:

The target system must be in a logically clean state. If any of the following components are available, the target is not considered clean and the command will fail:

- RAID groups
- LUNs
- User-created storage groups

Note

Duplication will not create the following logical components on the target system(s): SnapView sessions, SAN Copy information, synchronous/asynchronous mirror secondary information. If the XML captured from a mirror secondary system is used for duplication, then no mirror information will be created in the target system.

Note

While the `-capture` command supports both XML and CSV file formats, the `-duplicate` command processes information only in the XML format.

⚠ CAUTION

This command creates a new configuration in the system; use it only after completely understanding the command's behavior.

SYNTAX

```
arrayconfig -duplicate [-input path to local file name]
[-recover] [-stop] [-rollback] [-status] [-logfile] [-o]
```

OPTIONS

`-input path to local file name`

This is the filename of the captured information (captured from issuing the `configuration -capture` command).

`-logfile`

Retrieves and outputs the duplication log file stored on the targeted system.

`-o`

Executes the command without prompting for confirmation.

`-recover`

Used in a disaster recovery situation (if the system in trouble needs to be reconfigured from scratch using a configuration captured from the same system when it was healthy).

This assumes that the host connectivity and everything pertaining to the system will be the same. The initiators, the host objects and their associations with storage groups will be duplicated.

`-rollback`

Attempts to back out of the changes that are added to the configuration during the last duplication operation.

Changes made to the configuration after the duplication operation will not be accounted for, and may affect the rollback operation. For example, after a successful duplication, if you destroy a RAID group and create another RAID group with the same ID, the rollback will destroy the newly created RAID group as well. The rollback operates based on the object IDs stored in an internal context file. If the IDs are the same, it will not differentiate between the objects created by the previous duplication and the user-created objects.

`-status`

Provides a brief description of the current status of the duplication operation.

`-stop`

Terminates the currently running duplication operation, leaving the configuration intact at the point that the duplication ended.

Note

The `-logfile` and `-status` commands can be run with monitor privileges. All other `-duplicate` commands require manager privileges.

EXAMPLE

```
naviseccli -h ssl_spa arrayconfig -duplicate -input c:\Capture158.xml
```

arrayname

Renames a system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli arrayname` command changes the name of a system. If you omit a system name, the command returns the current system name (in the format host-device or system name).

The CLI prompts for confirmation as follows.

SYNTAX

```
arrayname NewArrayName
```

OPTIONS

```
arrayname NewArrayName
```

Specifies the new name for the system.

EXAMPLE

This command changes the name of the system to `users11`.

```
naviseccli -h ss1_spa arrayname users11
```

```
Change the arrayname from ss1 to users11? (y/n) y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not Supported` error message is printed to `stderr`.

backendbus -analyze

Displays speed and state information for each back-end bus and the devices on the back-end bus.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-analyze` function, displays the current and maximum back-end bus speeds for each back-end bus on the system, and each device (enclosures, disks, and link control cards (LCCs)) on the back-end bus. It also lists state information that includes whether a device on the bus is limiting and whether the system meets conditions required for a reset (see the list of conditions later in this section).

The current speed indicates the speed at which a bus or device is currently running; the maximum speed indicates the highest speed at which a bus or device is capable of running. If a device is limiting, it has a maximum speed lower than the maximum speed of the back-end bus. If you have a limiting device, a reset will not occur when you issue `backendbus-resetspeed`.

When you run `backendbus -analyze`, a summary analysis appears for each back-end bus to indicate whether a reset operation will be successful for the back-end bus. One of the following summary results will appear for each back-end bus.

Table 1 backendbus -analyze summary analysis

Summary analysis ^a	Description
Speed may change on rescan	A speed-related fault exists on the back-end bus (see Description section). The CLI cannot determine whether a reset will occur.
Speed will change on rescan	A reset will occur.
Speed will not change on rescan	There are limiting devices on the back-end bus. A reset will not occur.

- a. If a limiting device or speed-related fault exists on the back-end bus, a listing of the limiting devices and speed-related faults appears before the summary analysis. Devices with speed-related faults appear as Devices Requesting Bypass.

A speed-related fault can occur when you add a device to a back-end bus that has a maximum speed lower than the current speed of the bus. A device with a speed-related fault appears with unknown speed values; therefore, the CLI cannot determine if a reset will occur.

Note

If you issue the `backendbus -resetspeed` command, and a device with a speed-related fault exists on the bus, the reset could decrease the back-end bus speed (downgrade).

Your system must meet required conditions before you can issue a reset. The following table shows the list of conditions.

Table 2 Conditions

Condition	Description (pass criteria)
Analyze Speed Rule	The summary analysis for the back-end bus, displayed in <code>backendbus -analyze</code> , indicates that a reset will or may occur.
System Faults Rule	No system faults can exist.
Active Sessions Rule	No active SnapView, MirrorView, or SAN Copy sessions.
Transitions Rule	No active transitions such as, LUN binds, LUN expansions, RAID group expansions, disk equalization.
I/O Rate Rule	No active I/O.

Note

Run `backendbus -analyze` before you issue `backendbus -resetspeed`.

SYNTAX

```
backendbus -analyze
```

EXAMPLE

The following example displays current and maximum back-end bus speeds for each back-end bus on the specified system and the devices on the back-end bus:

```
naviseccli -h ss1_spa -user user11 -password pswd11 -scope 0
backendbus -analyze
```

OUTPUT

Bus 0 - Current Speed: 6Gbps. Maximum speed: 6Gbps.

Device	Current Speed	Max Speed	Limiting	Type
Bus 0 Enclosure 0	6	6	No	DAE6S
LCC A	6	6	No	
LCC B	6	6	No	
Disk 0	6	6	No	SAS
Disk 1	6	6	No	SAS
Disk 2	6	6	No	SAS
Disk 3	6	6	No	SAS
Disk 4	6	6	No	SAS
Disk 5	6	6	No	SAS
Disk 6	6	6	No	SAS
Disk 7	6	6	No	SAS
Disk 8	6	6	No	SAS
Disk 9	6	6	No	SAS
Disk 10	6	6	No	SAS
Disk 11	6	6	No	SAS
Disk 12	6	6	No	SAS
Disk 13	6	6	No	SAS
Disk 14	6	6	No	SAS
Bus 0 Enclosure 1	6	6	Maybe	DAE7S
ICM A	6	6	Maybe	
ICM B	6	6	Maybe	
LCC C	6	6	Maybe	
LCC D	6	6	Maybe	
Disk A0	6	6	No	SAS
Disk A1	6	3	No	SATA
Disk A2	6	6	No	SAS
Disk A3	6	3	No	SATA
Disk A4	6	3	No	SATA
Disk A5	6	3	No	SATA
Disk A6	Unknown	Unknown	No	Unknown
Disk A7	Unknown	Unknown	No	Unknown
Disk A8	Unknown	Unknown	No	Unknown
Disk A9	Unknown	Unknown	No	Unknown
Disk A10	Unknown	Unknown	No	Unknown
Disk A11	Unknown	Unknown	No	Unknown
Disk B0	6	3	No	SATA
Disk B1	6	6	No	SAS
Disk B2	6	3	No	SATA
Disk B3	6	3	No	SATA
Disk B4	6	3	No	SATA
Disk B5	6	3	No	SATA
Disk B6	Unknown	Unknown	No	Unknown
Disk B7	Unknown	Unknown	No	Unknown
Disk B8	Unknown	Unknown	No	Unknown
Disk B9	Unknown	Unknown	No	Unknown
Disk B10	Unknown	Unknown	No	Unknown
Disk B11	Unknown	Unknown	No	Unknown
Disk C0	6	3	No	SATA
Disk C1	6	3	No	SATA
Disk C2	6	3	No	SATA
Disk C3	6	3	No	SATA
Disk C4	6	3	No	SATA
Disk C5	6	3	No	SATA
Disk C6	Unknown	Unknown	No	Unknown

```

Disk C7           Unknown      Unknown      No           Unknown
Disk C8           Unknown      Unknown      No           Unknown
Disk C9           Unknown      Unknown      No           Unknown
Disk C10          Unknown      Unknown      No           Unknown
Disk C11          Unknown      Unknown      No           Unknown
Disk D0           6           6           No           SAS
Disk D1           6           6           No           SAS
Disk D2           6           3           No           SATA
Disk D3           6           3           No           SATA
Disk D4           6           3           No           SATA
Disk D5           6           3           No           SATA
Disk D6           Unknown      Unknown      No           Unknown
Disk D7           Unknown      Unknown      No           Unknown
Disk D8           Unknown      Unknown      No           Unknown
Disk D9           Unknown      Unknown      No           Unknown
Disk D10          Unknown      Unknown      No           Unknown
Disk D11          Unknown      Unknown      No           Unknown
Disk E0           6           3           No           SATA
Disk E1           6           3           No           SATA
Disk E2           6           3           No           SATA
Disk E3           6           3           No           SATA
Disk E4           6           3           No           SATA
Disk E5           6           3           No           SATA
Disk E6           Unknown      Unknown      No           Unknown
Disk E7           Unknown      Unknown      No           Unknown
Disk E8           Unknown      Unknown      No           Unknown
Disk E9           Unknown      Unknown      No           Unknown
Disk E10          Unknown      Unknown      No           Unknown
Disk E11          Unknown      Unknown      No           Unknown

```

```

-----
Bus 0 is currently running at 6Gbps.
Speed will not change on rescan

```

```

Analyze Speed Rule...Failed No loop will change speed on rescan.
System Faults Rule... Failed: Bus 0 Enclosure 1 Power A is Faulted
Active Sessions Rule... Passed
Transitions Rule... Passed
I/O Rate Rule... Passed

```

backendbus -get -connstate

Displays the connector state information.

PREREQUISITES

You must have a user account on the storage system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -connstate` switches displays the connector state for the mini-SAS HD connector(s) inserted into the BE SAS ports on the Mezzanine card.

SYNTAX

```
backendbus -get -connstate
```

EXAMPLE

The following output shows the output for the `backendbus -get -connstate` command when the mini SAS HD cable is connected on the Hypernova SLIC:

```

Bus 0
SPA Connector State: Online
SPB Connector State: Online

```

backendbus -get -physical

Displays physical slot and physical port ID information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -physical` switches displays the physical location of the back-end bus on the VNX series system. The physical location information includes I/O module and the physical port ID.

The `backendbus -get -all` command also displays the physical location information.

SYNTAX

```
backendbus -get -physical [backend bus number]
```

OPTIONS

backend bus number

Specifies the back-end bus for which you want to display bus information.

EXAMPLE

```
naviseccli -h ssl_spa backendbus -get -physical
```

OUTPUT

```
Bus 0
I/O Module Slot: 0
Physical Port ID: 0
Port Combination in Use: YES

Bus 0
I/O Module Slot: 0
Physical Port ID: 1
Port Combination in Use: YES

Bus 2
I/O Module Slot: 0
Physical Port ID: 2
Port Combination in Use: NO
```

backendbus -get -prominfo

Displays information returned from the resume prom in the connectors.

PREREQUISITES

You must have a user account on the storage system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -prominfo` switch displays the new information returned from the resume prom in the SFP+ and mini-SAS HD connectors.

SYNTAX

```
backendbus -get -prominfo
```

EXAMPLE

The following output shows the output for the `backendbus -get -prominfo` command when the mini-SAS HD cable is connected on the Hypernova SLIC:

```
Bus 0
SPA SFP/Connector EMC Part Number: TBD
SPB SFP/Connector EMC Part Number: TBD

SPA SFP/Connector EMC Serial Number: TBD
SPB SFP/Connector EMC Serial Number: TBD

SPA SFP/Connector Vendor Part Number: TBD
SPB SFP/Connector Vendor Part Number: TBD

SPA SFP/Connector Vendor Serial Number: TBD
SPB SFP/Connector Vendor Serial Number: TBD

SPA SFP/Connector Supported Speeds:
3Gbps.
6Gbps.
SPB SFP/Connector Supported Speeds:
3Gbps.
6Gbps.
```

backendbus -get -sfpstate

Displays the current condition of each back-end SFP.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -sfpstate` function displays the current condition of each back-end SFP, which is a small form-factor pluggable (SFP) optical modular transceiver. You can target a specific back-end bus and display information for only that bus, or you can display back-end bus information for all buses on the system.

The SFP states are Online, Faulted, and Removed. Online indicates that the back-end SFP is functioning normally; Faulted indicates an unsupported or faulted SFP; and Removed indicates that it does not exist.

Use `backendbus -analyze` to display speed information for each device on a back-end bus.

SYNTAX

```
backendbus -get -sfpstate [backend bus number]
```

OPTIONS

backend bus number

Specifies the back-end bus for which you want to display bus information.

EXAMPLE

The following example displays the current condition for each back-end SFP:

```
naviseccli -h ss1_spa backendbus -get -sfpstate
```

OUTPUT

```
Bus 0:
```

```
SPA SFP State: Faulted
SPB SFP State: Online
```

backendbus -get -speeds

Displays back-end bus speed information for the back-end buses on a system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -speeds` function displays the current and available back-end bus speeds, for the back-end buses on a system. You can target a specific back-end bus and display only information for that bus, or you can display back-end bus speed information for all buses on the system.

The current speed indicates the speed the bus is currently running at; the available speed(s) indicates the speed(s) available to the back-end bus. If the current speed is less than the maximum available speed, you can run the `backendbus -analyze` command to verify if there are limiting devices on the back-end bus. If a device is limiting, it has a maximum speed lower than the maximum speed of the back-end bus. If you have a limiting device, a reset will not occur when you issue `backendbus -resetspeed`.

Use `backendbus -analyze` to display speed information for each device on a back-end bus.

SYNTAX

```
backendbus -get -speeds [backend bus number]
```

OPTIONS

backend bus number

Specifies the back-end bus for which you want to display bus speed information.

EXAMPLE

The following example displays current and available bus speed information for back-end bus number 2:

```
naviseccli -h ss1_spa backendbus -get -speeds 2
```

OUTPUT

```
Bus 2:
Current Speed: 6Gbps.
Available Speeds:
3Gbps.
6Gbps.
```

backendbus -get -all

Displays the bus speeds, SFP state, and physical location of the bus.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-get -all` function displays the bus speeds and the current condition of each back-end SFP, which is a small form-factor pluggable (SFP) optical modular transceiver for the back-end buses on a system. It displays the physical location of the back-end bus (I/O module and the physical port ID)

only for the VNX series system. You can target a specific back-end bus and display information for only that bus, or you can display back-end bus information for all buses on the system.

The current speed indicates the speed the bus is currently running at; the available speed(s) indicates the speed(s) available to the back-end bus. If the current speed is less than the maximum available speed, you can run the `backendbus -analyze` command to verify if limiting devices exist on the back-end bus. If a device is limiting, it has a maximum speed lower than the maximum speed of the back-end bus. If you have a limiting device, a reset will not occur when you issue `backendbus -resetspeed`.

The SFP states are Online, Faulted, and Removed. Online indicates that the back-end SFP is functioning normally; Faulted indicates an unsupported or faulted SFP; and Removed indicates that it does not exist.

Use `backendbus -analyze` to display speed information for each device on a back-end bus.

SYNTAX

```
backendbus -get -all [backend bus number]
```

OPTIONS

backend bus number

Specifies the back-end bus for which you want to display bus information.

EXAMPLE

The following example displays the bus speed, current condition, and physical location for each back-end SFP:

```
naviseccli -h ss1_spa backendbus -get -all
```

OUTPUT

```
Bus 0

Current Speed: 6Gbps.
Available Speeds:
0Gbps.
3Gbps.
6Gbps.
SPA SFP State: N/A
SPB SFP State: N/A
I/O Module Slot: Base Module
Physical Port ID: 0
Port Combination In Use: No

SPA Connector State: None
SPB Connector State: None

SPA SFP/Connector EMC Part Number: N/A
SPB SFP/Connector EMC Part Number: N/A

SPA SFP/Connector EMC Serial Number: N/A
SPB SFP/Connector EMC Serial Number: N/A

SPA SFP/Connector Vendor Part Number: N/A
SPB SFP/Connector Vendor Part Number: N/A

SPA SFP/Connector Vendor Serial Number: N/A
SPB SFP/Connector Vendor Serial Number: N/A

SPA SFP/Connector Supported Speeds:
N/A
```

```

SPB SFP/Connector Supported Speeds:
N/A

Bus 1

Current Speed: 0Gbps.
Available Speeds:
0Gbps.
3Gbps.
6Gbps.
SPA SFP State: N/A
SPB SFP State: N/A
I/O Module Slot: Base Module
Physical Port ID: 1
Port Combination In Use: No

SPA Connector State: None
SPB Connector State: None

SPA SFP/Connector EMC Part Number: N/A
SPB SFP/Connector EMC Part Number: N/A

SPA SFP/Connector EMC Serial Number: N/A
SPB SFP/Connector EMC Serial Number: N/A

SPA SFP/Connector Vendor Part Number: N/A
SPB SFP/Connector Vendor Part Number: N/A

SPA SFP/Connector Vendor Serial Number: N/A
SPB SFP/Connector Vendor Serial Number: N/A

SPA SFP/Connector Supported Speeds:
N/A
SPB SFP/Connector Supported Speeds:
N/A

```

The following output shows the output for the `backendbus -get -all` command when the mini-SAS HD cable is connected on the Hypernova SLIC:

```

Bus 0

Current Speed: 6Gbps.
Available Speeds:
3Gbps.
6Gbps.

SPA SFP State: N/A
SPB SFP State: N/A

I/O Module Slot: 1
Physical Port ID: 0

SPA Connector State: Online
SPB Connector State: Online

SPA SFP/Connector EMC Part Number:          TBD
SPB SFP/Connector EMC Part Number:          TBD

SPA SFP/Connector EMC Serial Number:        TBD
SPB SFP/Connector EMC Serial Number:        TBD

SPA SFP/Connector Vendor Part Number:       TBD
SPB SFP/Connector Vendor Part Number:       TBD

```



```

SPA SFP/Connector Vendor Serial Number:      TBD
SPB SFP/Connector Vendor Serial Number:      TBD

SPA SFP/Connector Supported Speeds:
3Gbps.
6Gbps.
SPB SFP/Connector Supported Speeds:
3Gbps.
6Gbps.

```

The following example shows the output for the `backendbus -get -all` command when the mini-SAS HD cable is connected:

```

Bus 0

Current Speed: 6Gbps.
Available Speeds:
3Gbps.
6Gbps.

SPA SFP State: N/A
SPB SFP State: N/A

I/O Module Slot: Onboard
Physical Port ID: 0

SPA Connector State: Online
SPB Connector State: Online

SPA SFP/Connector EMC Part Number: N/A
SPB SFP/Connector EMC Part Number: N/A

SPA SFP/Connector EMC Serial Number: N/A
SPB SFP/Connector EMC Serial Number: N/A

SPA SFP/Connector Vendor Part Number: N/A
SPB SFP/Connector Vendor Part Number: N/A

SPA SFP/Connector Vendor Serial Number: N/A
SPB SFP/Connector Vendor Serial Number: N/A

SPA SFP/Connector Supported Speeds:
N/A
SPB SFP/Connector Supported Speeds:
N/A

```

backendbus -resetspeed

Resets the back-end bus speeds for a system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli backendbus` command with the `-resetspeed` function, resets the back-end bus speed for the buses on the system. When you execute `backendbus -resetspeed`, if the back-end bus is eligible for a reset and the system meets the conditions required for a back-end bus speed reset, the reset occurs. See the table below for the steps you should take before issuing `backendbus -resetspeed`.

Note

This is a disruptive operation and can take several minutes to complete. During this time the system will be offline and no I/O will be performed.

Before you execute a reset, run `backendbus -analyze` to verify the current and maximum speeds of each back-end bus and device to determine if limiting devices exist, which will prevent a reset, and to verify that your system meets the conditions required for a reset.

To verify whether the back-end bus speed reset occurs, capture the output from `backendbus -analyze` or `backendbus -get -speeds` before you execute the reset, to compare to the reset back-end bus speed values after you issue the reset.

Steps to take before issuing a reset:

1. Run `backendbus -analyze` to verify the following:
 - back-end bus and device speed information
 - no limiting devices exist on the back-end bus
 - system meets required conditions
2. Save the output from `backendbus -analyze` or `backendbus -get -speeds`, to compare the bus speed in the saved output to the bus speed after you initiate `backendbus -resetspeed`.

SYNTAX

```
backendbus -resetspeed
```

EXAMPLE

The following example resets the the back-end bus speeds for the specified system:

```
naviseccli -h ssl_spa backendbus -resetspeed
```

```
WARNING: the command you are about to execute requires
that the system be completely shutdown. During this period
of time all IO to attached application hosts will be
suspended. All affected applications must be quiesced
before proceeding to avoid data unavailability and
possible data loss. This operation will take several
minutes to complete.
```

```
Are you sure you want to proceed? (y/n)?
```

OUTPUT

None if the command succeeds; status or error information if it fails.

baseuid

Gets or sets the system unique unit identifier (UUID) base for use by a server running Tru64.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli baseuid` command gets or sets the base of the unique unit ID as required for servers in a Tru64 cluster. The Tru64 operating system uses the UUID (which it creates by combining the base and an offset, which is the system LUN ID) to manage the LUNs in a cluster. You can also use the UUID to manage the LUN.

To discover the UUID base, use the command without an argument.

SYNTAX

`baseuuid [uuid-base-value]`

OPTIONS

uuid-base-value
Specifies the UUID base.

EXAMPLE

These commands get, then set the UUID base of the system containing the SP `ss1_spa`.

`naviseccli -h ss1_spa baseuuid`

Base UUID of the system: 12345

`naviseccli -h ss1_spa baseuuid 34`

Base UUID of the system is set to 34.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not Supported` error message is printed to `stderr`.

bind

Creates a LUN within an existing RAID group.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli bind` command binds a LUN within an existing RAID group. You can create a RAID group with the `createrg` command (see [createrg on page 56](#)).

Ensure that you address the `bind` command to the SP that you want to own the LUN.

The following table identifies the parameters you must configure for each RAID type you bind.

Table 3 LUN parameters for different RAID types

RAID type	Bind parameters
RAID 6, RAID 5, RAID 1/0	LUN number Write caching
RAID 3	LUN number
RAID 1	LUN number Write caching
RAID 0	LUN number Write caching
Disk (individual disk unit)	LUN number Write caching

Table 3 LUN parameters for different RAID types (continued)

To allow for good choices of available disk modules, bind LUNs in this order:

1. First - RAID 1/0
2. Second - RAID 6, RAID 5, RAID 3, and RAID 0
3. Third - RAID 1
4. Fourth - Individual units

Before you start binding disk modules, read the restrictions and recommendations in the following table.

Table 4 Restrictions and recommendations for binding disks into LUNs

LUN to bind	Restrictions and recommendations
Any LUN	You can bind only unbound disk modules. All disk modules in a LUN must have the same capacity to fully use the modules' storage space.
RAID 6	A RAID 6 group has an even number of disks. You must bind a minimum of 4 disk modules and no more than 16 disk modules. We recommend that you bind 6 or 12 modules for more efficient use of disk space. Write caching is highly recommended for RAID 6.
RAID 5	You must bind a minimum of 3 disk modules and no more than 16 disk modules. We recommend you bind 5 modules for more efficient use of disk space.
RAID 3	You must bind exactly 5 or 9 disk modules. Select 5 drives for quicker rebuild times or 9 drives for more efficient use of space.
RAID 1	You must bind exactly 2 disk modules.
RAID 0	You must bind a minimum of 3 disk modules, and no more than 16 disk modules.
RAID 1/0	You must bind a minimum of 2 disk modules, and an even number of modules, but no more than 16 modules. Unisphere pairs modules into mirrored images in the order in which you select them. The first and second modules you select are a pair of mirrored images; the third and fourth modules you select are another pair of mirrored images; and so on. The first module you select in each pair is the primary image, and the second module is the secondary image.
Individual disk unit	None

Whether you should use caching for a specific LUN depends on the RAID type. Recommendations for caching with different RAID types follow.

Table 5 Caching recommendations for different RAID types

RAID type	Write caching
RAID 0	Acceptable
RAID 1	Acceptable
RAID 1/0	Acceptable

Table 5 Caching recommendations for different RAID types (continued)

RAID type	Write caching
RAID 3	Recommended
RAID 5	Highly recommended
RAID 6	Highly recommended
Disk	Acceptable

Note

When binding RAID 1/0 LUNs, the order in which you select modules is important. The first module you select is the first part of the mirror's primary image and the second module is the first part of the mirror's secondary image; the third module is the second part of the primary image, the fourth module is the second part of the secondary image; and so on for any other modules in the LUN. In other words, the first and second modules selected are a pair of peer image parts, the third and fourth modules are another pair of peer image parts, and so on.

If you need to change the description (RAID type), you must unbind the LUN ([unbind on page 181](#)), and then rebind it with the correct parameters. If you need to change the write cache state, use the `chglun` command ([chglun on page 50](#)).

SYNTAX

```
bind [raidtype] [lun] -rg rgID [-cnt count_luns_to_create]
[-aa auto_assignment] [-cap capacity] [-fastcache 0|1]
[-noinitialverify] [-pl placement]
[-sp a|b] [-sq size-qualifier]
[-wc writecache]
```

OPTIONS*raid-type*

Specifies the RAID type for the LUN.

r0

= RAID 0

r1

= RAID 1

r3

= RAID 3

r5

= RAID 5

r6

= RAID 6

r1_0

= RAID 1/0

id
= individual disk

Note

All RAID types are not supported on all systems.

lun

Specifies a decimal number to assign to the LUN valid range as shown in *VNX for Block Concepts* guide, in section LUN IDs, unique IDs and disk IDs. If a valid range is not specified, the lowest available number (the default value) is assigned to the LUN.

-rg rgID

The RAID group identification number. The RAID group must already exist. You can create a RAID group with the `createrg` command ([createrg on page 56](#)).

rgID specifies a number between 0 and the maximum number of RAID groups supported by the RAID group system.

-cnt

Specifies the number of LUNs to create.

-aa auto_assignment (CX series only)

Enables or disables autoassignment functionality for the LUN. *auto_assignment* specifies functionality as follows:

0

Disables auto-assignment

1

Enables auto-assignment (the default)

-cap capacity

Sets the capacity of usable space in the LUN. (The default is full space available.)

-fastcache 0|1

Enables or disables FAST Cache on the LUN.

-noinitialverify

Does not perform an initial background verification. When this switch is not present, the bind performs an initial background verification on the newly bound LUN and eliminates any latent soft media errors.

-pl placement

Places a LUN in a RAID group; placement specifies which algorithm to use as follows:

bf

best fit algorithm (the default)

ff

first fit algorithm

-sp a|b

Sets the default owner of the LUN: SP A or SP B. The default is the current SP.

-sq size-qualifier

Sets the default size qualifier. *size-qualifier* specifies the qualifier as follows:

mb

megabytes

gb

gigabytes (the default)

tb
terabytes

bc
block count

-wc *write-cache*

Enables or disables the system write cache for this specific LUN. *writecache* values:

0
Disables write cache

1
Enables write cache (the default).

EXAMPLE

```
naviseccli bind 10 -rg 0 -sq gb -cap 1
```

OUTPUT

There is no output. Errors are printed to `stderr`.

cachecard -list

Lists the state information of the cache memory card.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cachecard` command with the `-list` function lists the state information of the cache memory card. The memory size and hardware state of the cache memory card is displayed.

You can choose to display only the memory size by specifying the `-size` switch or display only the hardware state by specifying the `-hwstate` switch. Without these switches, the memory size and hardware state are displayed.

The possible hardware states are listed in the following table:

Table 6 Hardware states

Hardware state	Meaning
OK	The power on the main SP is on and the cache memory card battery is fully charged. The cache memory card is in normal operation.
Not present	The cache memory card is not present.
Faulted	The cache memory card has hardware faults.
Battery charging	The power on the main SP is on and the battery is charging on the cache memory card.
In wrong array	The cache memory card is from the wrong system.

SYNTAX

```
cachecard -list [-size] [-hwstate]
```

OPTIONS`-size`

Specifies that only the total memory size of the cache memory card (MB) be displayed.

`-hwstate`

Specifies that only the hardware state of the cache memory card be displayed.

EXAMPLE

This command displays all state information (memory size and hardware state) of the cache memory card.

```
naviseccli -h ssl_spa cachecard -list
```

OUTPUT

```
Total Memory:    512MB
Hardware State:  Ok
```

cachecard -initialize

Initializes the cache memory card.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cachecard` command with the `-initialize` function lets you initialize the cache memory card.

The `cachecard -initialize` command can be issued under any of the following circumstances:

- a different cache memory card is installed on the system
- live data is on the cache memory card
- a LUN with an offline (cache dirty) condition exists on the system

SYNTAX

```
cachecard -initialize [-o]
```

OPTIONS`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For this example, a different cache memory card is installed on the system.

```
naviseccli -h ssl_spa cachecard -initialize
```

```
The cache memory card currently installed in this system is not the
same as the one previously installed in
this system. In order for this system to
function properly, the current cache memory card must be
initialized on this system.
```

```
Do you want to initialize Cache Memory Card now? (y/n)?
```

OUTPUT

None if the command succeeds; status or error information if it fails.

cache -sp -info

Lists the SP Cache properties.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cache -sp -info` command lists all SP cache properties, state information, and performance statistics.

SYNTAX

```
cache -sp -info [-state] [-perfdata]
```

OPTIONS

`-state`

Displays SP Cache state, which may be one of the following:

- Enabled
- Disabled
- Disabling
- Enabled with dirty data for unwritable LUNs
- Disabling with dirty data for unwritable LUNs

`-perfdata`

Displays SP Cache performance statistics.

EXAMPLE #1

This example displays SP Cache state.

```
naviseccli -h ssl_spa cache -sp -info -state
```

```
SP Read Cache State: Enabled
SP Write Cache State: Enabled
```

EXAMPLE #2

This example displays SP Cache performance statistics.

```
naviseccli -h ssl_spa cache -sp -info -perfdata
```

```
Cache Page size (KB) : 8
Read Hit Ratio       : 100
Write Hit Ratio      : 100
Dirty Cache Pages (MB) : 0
```

EXAMPLE #3

This example displays SP Cache properties and statistics.

```
naviseccli -h ssl_spa cache -sp -info
```

```
Cache Information
-----
SP Read Cache State: Enabled
SP Write Cache State: Enabled
Cache Page size (KB): 8
Read Hit Ratio: 100
Write Hit Ratio: 100
Dirty Cache Pages (MB): 0
```

```
SPA Read Cache State: Enabled
SPB Read Cache State: Enabled
SPA Write Cache State: Enabled
SPB Write Cache State: Enabled
SPS Test Day: Saturday
SPS Test Time: 20:00
SPA Physical Memory Size (MB): 65536
SPB Physical Memory Size (MB): 65536
```

cache -sp -modify

Modifies the SP Cache properties.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cache -sp -modify` command modifies the SP cache properties.

SYNTAX

```
cache -sp -modify-wc [on|off] [-o]
```

OPTIONS

`-wc [on | off]`

Enables or disables write cache.

`on` Enable write cache

`off` Disable write cache

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This example disables write cache.

```
naviseccli -h ss1_spa cache -sp -modify -wc off
```

```
Warning! Disabling write cache at the system level may seriously
degrade system performance and is not recommended.
Continue disabling system write cache? (y/n)?
```

OUTPUT

There is no output. Errors are printed to `stderr`.

chglun

Changes parameters on a currently bound LUN.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli chglun` command changes parameters on a currently bound LUN. You must use at least two switches with `chglun` and one of the switches must be the LUN number `-l lun`.

Note

The `chglun` command does not support thin LUNs.

To change a metaLUN, use the `naviseccli` command with the `-modify` switch, as described in [metalun -modify on page 196](#).

Table 7 Caching recommendations for different RAID types

RAID type	Write caching
RAID 0	Acceptable
RAID 1	Acceptable
RAID 1/0	Acceptable
RAID 3	Recommended
RAID 5	Highly recommended
RAID 6	Highly recommended
Disk	Acceptable

Note

When binding RAID 1/0 LUNs, the order in which you select modules is important. The first module you select is the first part of the mirror's primary image and the second module is the first part of the mirror's secondary image; the third module is the second part of the primary image, the fourth module is the second part of the secondary image; and so on for any other modules in the LUN. In other words, the first and second modules selected are a pair of peer image parts, the third and fourth modules are another pair of peer image parts, and so on. Any change to the rebuild time takes effect as soon as the operation completes.

SYNTAX

```
chglun -l lun [-a autoassign] [-c cachetype] [-d defaultowner]
[-fastcache 0|1] [-namenewname] [-w writeaside]
```

OPTIONS

`-l lun`

Specifies which LUN to change. *lun* specifies the LUN number. This switch is required.

`-a auto-assign`

Enables or disables automatic reassignment of LUN ownership to an active SP if an SP failure occurs.

auto_assign specifies functionality as follows:

0

Disables autoassignment for specified LUN

1

Enables autoassignment for specified LUN

`-c cache-type`

Enables read caching, or both read and write caching. For caching to occur, system caching must also be enabled. Whether or not you should use caching for a specific LUN depends on the RAID type. See the [caching recommendations for different RAID types on page 51](#).

cache-type specifies the type of caching used as follows:

read

read caching

rw

read and write caching

-d *default-owner*

Changes the default owner of the LUN. *default-owner* specifies ownership as follows:

0

Changes ownership of the LUN to SP A

1

Changes ownership of the LUN to SP B

-fastcache 0|1

Modifies the FAST Cache configuration of the LUN.

-name *newname*

Sets the LUN name. *newname* specifies the new name for the LUN. This switch applies only to systems running a version of VNX OE for Block software that supports storage group commands. The LUN name is displayed in the UI and has no other purpose. You cannot use the LUN name as a CLI command argument. You can assign the same name to multiple LUNs.

Note

The **-name** switch applies only to shared systems.

-w *write-aside* (Dual-SP AX4-5 series running Unisphere and CX series only)

write-aside specifies, in blocks, the largest write request size that will be written to cache. Write requests greater than the *write-aside* value are written directly to disk, bypassing write cache. Valid values are 0 through 65536. The default value '0' indicates that all writes will be cached.

EXAMPLE

Change LUN 3 to enable read and write caching.

```
naviseccli -h ss1_spa chglun -l 3 -c rw
```

OUTPUT

There is no output. Errors are printed to `stderr`.

chgrg

Changes properties of a RAID group.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli chgrg` command changes the properties of a specified RAID group. You must use only one of the optional switches with `chgrg` and must specify the RAID group identification number, *rgID*.

Note

The `chgrg` command does not support private RAID groups used in thin pools.

SYNTAX

```
chgrg rgID
[-powersavings on|off]
```

OPTIONS

rgID

Is the RAID group identification number. *rgID* specifies a number between 0 and the maximum number of RAID groups supported by the RAID group system.

`-powersavings on|off`

Configures the power savings settings. If the selected disks are not eligible for power savings settings, the system displays an error message.

EXAMPLE

This command turns the power savings settings ON in RAID group 2.

```
naviseccli -h ss1_spa chgrg 2 -powersavings on
```

OUTPUT

There is no output. Errors are printed to `stderr`.

clearlog

Clears the SP's error log.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli clearlog` command deletes the contents of the SP's unsolicited error log.

The command clears only the log of the SP you specify with the `-h` switch in the command.

SYNTAX

```
clearlog
```

EXAMPLE

This command deletes the SP logs of the system.

```
naviseccli -h ss1_spa clearlog
```

```
unsolicited error log cleared
```

OUTPUT

Confirmation message.

clearstats

Resets the SP's statistics logging.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli clearstats` command resets statistics logging on the SP. If statistics logging is on, it is turned off, and then back on. If statistics logging is off, it is turned on, and then off.

The command affects only the log of the SP you specify with the `-h` switch in the command.

SYNTAX

```
clearstats
```

EXAMPLE

```
naviseccli -h ssl_spa clearstats
```

```
Statistics successfully cleared, statistics logging disabled.
```

OUTPUT

Confirmation message.

cmdtime

Displays the date and endtime of command execution.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

This is a global command switch used to display the date and time at the end of each command execution. The time is displayed at the end of command output. This is an optional switch supported on all CLI commands. This switch cannot work independently.

`cmdtime` switch is used with the CLI commands. You must add the `-cmdtime` switch for each command to get the time information.

SYNTAX

```
cmdtime
```

OPTIONS

```
-cmdtime
```

Displays the date and end time of command execution.

EXAMPLE

```
naviseccli -h ssl_spa -cmdtime getagent
```

```
Agent Rev:          6.28.0 (1.83)
Name:               K10
Desc:
Node:               A-CF2VH072500004
Physical Node:     K10
```

```
Signature:          1962014
Peer Signature:    0
Revision:          4.28.40.1.093
SCSI Id:           0
Model:             CX4-40
Model Type:        Rackmount
Prom Rev:          0.47.00
SP Memory:         5120
Serial No:         CF2VH072500004
SP Identifier:     A
Cabinet:           SPE
CmdTime: 10/11/07 07:46:44
```

copytodisk

Copies data from a configured drive (part of a RAID Group) to an unbound drive.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli copytodisk` command initiates the copying of data from a configured drive (part of a RAID Group) to an unbound drive. The user can use this command to copy data from any bound disk to any unbound disk, not just from a permanent spare to a replacement drive.

SYNTAX

```
copytodisk sourcedisk [destdisk] [-o]
```

OPTIONS

`sourcedisk`

Specifies the source disk.

`destdisk`

Specifies the destination disk.

Note

The destination disk must be the same type as the source disk, and the same size or larger than the source disk. If the destination disk is not specified, the system will choose an appropriate disk to copy to.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE #1

This example initiates the copying of data from a bound disk `0_1_5` to an unbound disk `0_1_6`.

```
naviseccli copytodisk 0_1_5 0_1_6
```

```
WARNING: The data from disk 0_1_5 will be copied to disk 0_1_6. This
process cannot be aborted and may take a long time to complete.
Also current selection of disks will break the hot spare policy;
would you like to continue the copy? (y/n) y
```

EXAMPLE #2

This example initiates the copying of data from disk `0_1_5` to disk `0_1_6` which is part of an active `copytodisk` operation or is a bound drive.

```
naviseccli copytodisk 0_1_5 0_1_6
```

```
WARNING: The data from disk 0_1_5 will be copied to disk 0_1_6. This
process cannot be aborted and may take a long time to complete.
Do you wish to continue? (y/n) y
```

```
Error: copytodisk command failed
Error returned from Agent
SP A: The destination drive is busy processing a copy or is being
used by another RAID group. Please wait for the drive operation to
complete or copy to another drive. (0x717f8007)
```

EXAMPLE #3

This example initiates a proactive copy. A user may proactively initiate a copy and allow the system to pick the drive that it will use to copy the data to. This might be done if the user saw errors on the disk and wanted to copy the data before the drive actually failed. For VNX, the source disk is not taken off-line after the copy. A user could also initiate a proactive copy and supply the destination drive.

```
naviseccli copytodisk 0_0_7
```

```
WARNING: The disk 0_0_7 is selected for hot sparing. Its serial
number is 6SE4LHZE.
The proactive sparing process cannot be aborted and may take a long
time to complete. This disk will still be available after the
proactive copy is complete, and its state will change to Unbound.
Do you wish to continue? (y/n) y
```

createrg

Creates a RAID group.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli createrg` command creates a RAID group on the specified disks.

To create RAID groups on a systems running Navisphere Express, you must define the `-raidtype` switch.

Note

You cannot create private RAID groups using the `createrg` command.

SYNTAX

```
createrg rgID disks
-raidtype type [-powersavings on|off]
[-elementsize elementsize] [-o]
```

OPTIONS

rgID

Is the RAID group identification number. *rgID* specifies a number between 0 and the maximum number of RAID groups supported by the RAID group system.

disks

Specifies the physical disks included in this RAID group. In a RAID 1/0 or RAID 1 configuration, you can manually select the disks for the RAID group in an order that determines the primary and secondary image. For example, for a 1/0 RAID group, the

first two disks you list in the command form a pair of mirrored images, and the next two disks form another pair of mirrored images, and so on.

Refer to the *VNX for Block Concepts* guide for more information about disk IDs to verify the format for specifying disks.

Note

The disks in a RAID group must be SAS, NL-SAS, or different types of Flash drives. In general, you must use the same type of drive in the RAID Group for all the disks.

`-raidtype type`

Sets the RAID type of the RAID group during the time of creation. The supported RAID types are r6, r5, r3, r1_0, r0 and disk.

NOTICE

The `-raidtype` is a mandatory switch for VNX series systems.

`-powersavings on|off`

Turns the power savings settings ON or OFF on a RAID group. If `-powersavings` option is not specified, then RAID group is created with power savings off by default. If you specify disks that are not eligible for power savings settings, the system displays an error message.

`-elementsize elementsize`

Enables the `createrg` command to accept an element size for RAID groups. However, only the 4+1 RAID5 unit type, on non-Flash drives, accepts an element size as large as 1024 sectors or blocks. If a size of 1024 is requested for RAID groups other than 4+1 RAID5, the system displays the following error message:

```
Invalid RAID type for element size requested
Valid RAID types: r5 (with five disks)
```

If Flash drives are chosen for the RAID Group, with an element size of 1024, the system displays the following error message:

```
Invalid drive type for element size requested
Valid drive types: non-Flash drive types
```

Note

If no element size is specified, an element size of 128 sectors or blocks will be considered as the default value.

`-o`

Executes the command without prompting for confirmation. Skips hotspare policy checking.

EXAMPLE

This command creates RAID group 0 from the five specified disks and sets `raidtype r5` as the RAID group.

```
naviseccli -h ss1_spa createrg 0 0_1 0_2 0_3 0_4 0_5 -raidtype r5
```

emconfiguration

Configures event monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli emconfiguration` command can be used to disable event monitoring of an SP, for a specified amount of time. It will then re-enable the response after the specified time is up.

SYNTAX

```
emconfiguration [-reloadconfig] [-disableresponse time]
[-enableresponse]
```

OPTIONS

`-disableresponse time`

This switch disables event monitoring of the SP for the amount of time specified.

Time is indicated in minutes, with a maximum of four hours.

`-enableresponse`

This switch re-enables the response.

`-reloadconfig`

This switch will reload a template file.

EXAMPLE

```
naviseccli -h ssl_spa emconfiguration -reloadconfig -disableresponse
time -enableresponse
```

environment

Reports inlet air temperature and input power.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `environment` command displays the inlet air temperature and input power of the system, disk-array enclosures, and standby power supply through the CLI and Unisphere GUI.

SYNTAX

```
environment -list -array [-power] [-present][[-average]]
-enclosure [bus_enc|spe] [-power|-intemp [f|c]]
[-present][[-average]]
-sps [-sp a|b] [-power] [-present][[-average]]
-all
```

OPTIONS

`-array`

Displays the present or average input power information on the system.

`-power`

Displays the input power information.

- present
Displays the current value, which is a sum of the present input power of all supported systems. The current value is computed as the 30-second average of the power consumption sampled every 3 seconds.
- average
Displays the average value. The average value is computed as the 60-minute rolling average of the present power consumption values. It requires an hour to calculate the correct value. N/A is displayed if duration of data is less than an hour.
- enclosure
Displays the present or average inlet air temperature or input power information on a specified disk-array enclosure.
- bus_enc|spe
Specifies either the bus enclosure or SPE enclosure.
- intemp f|c
Displays the inlet air temperature information. The f flag indicates Fahrenheit and the c flag indicates Celsius. The default value is c.
- sps
Displays the present or average input power information on a specified standby power supply.
- sp a|b
Specifies the standby power supply ID on which to display the information.
- all
Displays the following:
 - System input power information
 - Data Mover enclosure inlet air temperatures and input power information
 - Array input power information
 - Disk-array enclosure inlet air temperatures and input power information
 - Storage processor enclosure inlet air temperatures and input power information
 - Standby power supply input power information

EXAMPLE #1

The following example shows the output for the `environment -list -all` command for the DPE9

```
environment -list -all
```

```
Array
Input Power
Status: Valid
Present(watts): 180
Rolling Average(watts): 180
DPE9 Bus 0 Enclosure 0
Input Power
Status: Valid
Present(watts): 180
Rolling Average(watts): 180

Air Inlet Temperature
Status: Valid
Present(degree C): 28
Rolling Average(degree C): 28
Bus 0 Enclosure 0 BBU A
Input Power
Status: Invalid
Present(watts): Not Available
Rolling Average(watts): Not Available
Bus 0 Enclosure 0 BBU B
Input Power
```

```
Status: Invalid
Present(watts): Not Available
Rolling Average(watts): Not Available
```

EXAMPLE #2

The following example shows the output for the `environment -list -all` command for the SPE9.

```
environment -list -all
```

```
Array
Input Power
Status: Valid
Present(watts): 114
Rolling Average(watts): 114
SPE9 Enclosure SPE
Input Power
Status: Valid
Present(watts): 114
Rolling Average(watts): 114
Air Inlet Temperature
Status: Valid
Present(degree C): 24
Rolling Average(degree C): 24
DAE5S Bus 0 Enclosure 0
Input Power
Status: Valid
Present(watts): 163
Rolling Average(watts): 163
Air Inlet Temperature
Status: Valid
Present(degree C): 36
Rolling Average(degree C): 36
Enclosure SPE SPS Module A
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Enclosure SPE SPS Battery A
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Enclosure SPE SPS Module B
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Enclosure SPE SPS Battery B
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Bus 0 Enclosure 0 SPS Module A
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Bus 0 Enclosure 0 SPS Battery A
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
Bus 0 Enclosure 0 SPS Module B
Input Power
Status: Valid
```

```
Present(watts): 0
Rolling Average(watts): 0
Bus 0 Enclosure 0 SPS Battery B
Input Power
Status: Valid
Present(watts): 0
Rolling Average(watts): 0
```

failovermode

Changes the type of trespass needed for failover software (non-storage group only).

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli failovermode` command enables or disables the type of trespass needed for failover software. Use this command with certain types of failover software as suggested in *What are the Initiator, Arraycommpath, and Failovermode settings for PowerPath, DMP, PVLlinks, and native failover software*, which is on EMC Online Support.

The CLI may prompt for confirmation of some changes you specify unless you use the `-o` (override). To change the setting, enter `y`.

This command works for systems without storage groups only. For systems with storage groups, use the command `storagegroup sethost -failovermode` (see [storagegroup on page 168](#)).

Changing the failover mode setting will ask for confirmation unless you include the `-o` switch.

CAUTION

Do not issue this command unless you know the appropriate value for the type of failover software running on the attached servers. The command is designed for transition from one type of failover software to another. Changing the failover mode to the wrong value will make the system inaccessible.

To discover the current setting of `failovermode`, use the command without an argument.

SYNTAX

```
failovermode [0|1|2|3|4] [-o]
```

OPTIONS

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For `ss1_spa`, this command enables failover mode 0.

```
naviseccli -h ss1_spa failovermode 0
```

```
WARNING: Previous Failovermode setting will be lost!
DO YOU WISH TO CONTINUE? (y/n) y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
Command not supported Valid values are 0 to 3
```

faults -list

Lists status information for faulted components on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli faults` command with the `-list` function, provides status information for faulted hardware components, such as enclosures, disks, power supplies, storage processors, standby power supplies, as well as MirrorView/Asynchronous replication software sessions.

SYNTAX

```
faults -list
```

EXAMPLE

The following example lists the faulted hardware components and MirrorView software sessions on the system:

```
naviseccli -h ssl_spa faults -list
```

OUTPUT

```
Faulted Subsystem:  APM00043806491
Enclosure SPE :  Faulted
Enclosure SPE Power B :  Faulted
Bus 0 Enclosure 0 :  Faulted
Bus 0 Enclosure 0 Disk 7 :  Removed
Bus 0 Enclosure 0 Disk 8 :  Removed
Bus 0 Enclosure 0 Disk 9 :  Removed
Bus 0 Enclosure 0 Power B :  Faulted

FAR Mirror Faults Status Report
50:06:01:60:90:60:00:14:09:00:00:00:00:00:00:00:00:00:00:00:
Asynchronous mirror is faulted

50:06:01:60:90:60:00:14:06:00:00:00:00:00:00:00:00:00:0050:06:01:
60:90:60:03:99:Asynchronous mirror image is fractured

50:06:01:60:90:60:00:14:0B:00:00:00:00:00:00:00:00:00:
Asynchronous mirror is faulted

50:06:01:60:90:60:00:14:07:00:00:00:00:00:00:00:00:00:
Asynchronous mirror is faulted

50:06:01:60:90:60:00:14:08:00:00:00:00:00:00:00:00:00:0050:06:01:
60:90:60:03:99:Asynchronous mirror image is fractured

50:06:01:60:90:60:00:14:01:00:00:00:00:00:00:00:00:00:
Asynchronous mirror is faulted

50:06:01:60:90:60:00:14:11:00:00:00:00:00:00:00:00:00:
Asynchronous mirror is faulted

50:06:01:60:90:60:00:14:12:00:00:00:00:00:00:00:00:00:0050:06:01:
```

```
60:90:60:03:99:Asynchronous mirror image is fractured
```

flash -ioport

Flashes the LEDs on I/O ports.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli flash -ioport` command flashes the I/O ports on VNX systems. Flashes only the targeted SP owned I/O ports.

SYNTAX

```
flash -ioport [-fe logicalportID | -be logicalportID | [-pslotid
"onboard" | "basemodule" | physicalslotID -pport physicalportID] [-byusage [mirrorview]] on|off
```

OPTIONS

`-ioport`

Controls flashing of I/O ports on the target SP I/O modules.

`-fe logicalportID`

Turns flashing on/off for the front-end port with the logical ID specified on the target SP.

`-be logicalportID`

Turns flashing on/off for the back-end port with the logical ID specified on the target SP.

`-pslotid "onboard" | physicalslotID -pport physicalportID`

Turns flashing on/off for an physical slot/port located on the targeted SP's I/O module. The physical slot ID refers to the I/O module slot.

Note

`"basemodule"` is not supported on VNX8000 Block systems.

`-byusage mirrorview`

Turns flashing on/off for an I/O port depending on how the I/O port is being used. The only option is `mirrorview`.

`on|off`

Indicates whether flashing is turned on or off.

EXAMPLE #1

This example turns flashing on for a specified physical port on the onboard I/O module:

```
naviseccli -h 10.244.211.36 flash -ioport -pslotid onboard -pportid 3 on
```

```
Slot Onboard, Physical Port 3 (Uninitialized): ON
```

EXAMPLE #2

This example turns flashing on by specifying a logical port ID on the onboard I/O module:

```
naviseccli -h 10.244.211.36 flash -ioport -be 3 on
```

```
Slot Onboard, Physical Port 1 (BE 3): ON
```

EXAMPLE #3

This example turns flashing on by specifying the use of an I/O port on the onboard I/O module:

```
naviseccli -h 10.244.211.36 flash -ioport -fe 1 on
```

```
Slot Onboard, Physical Port 3 (FE 1): ON
```

EXAMPLE #4

This example turns flashing on by specifying the use of an I/O port on the removable I/O modules:

```
naviseccli -h 10.244.211.36 flash -ioport -fe 4 off
```

```
Slot 0, Physical Port 0 (FE 4): OFF
```

EXAMPLE #5

This example turns flashing on for a specified physical port on the removable I/O modules:

```
naviseccli -h 10.244.211.36 flash -ioport -pslotid 1 -pportid 0 on
```

```
Slot 1, Physical Port 0 (FE 8): ON
```

EXAMPLE #6

This example turns flashing on for a specified physical port on the base module:

```
naviseccli -h 10.244.211.36 flash -ioport -pslotid basemodule -pportid 0 on
```

```
Slot basemodule, Physical Port 0 (BE 0):ON
```

flash -iomodule

Turns the LEDs for the I/O ports on an I/O module on or off.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli flash -iomodule` command turns the flashing LEDs on/off for all the I/O ports on a specified module on the VNX system.

SYNTAX

```
flash -iomodule ["onboard"|"basemodule"|slotid] [on|off]
```

OPTIONS

`"onboard"|"basemodule"|slotid`

If the slotID is not specified, an error message is displayed.

Note

`"basemodule"` is not supported on VNX8000 Block systems.

on|off

Indicates whether flashing is turned on or off.

EXAMPLE #1

```
naviseccli -h 10.244.211.36 flash -iomodule onboard on
```

```
I/O Module Onboard: ON
```

EXAMPLE #2

```
naviseccli -h 10.244.211.36 flash -iomodule onboard off
```

```
I/O Module Onboard: OFF
```

EXAMPLE #3

```
naviseccli -h 10.244.211.36 flash -iomodule basemodule 0 on
```

```
I/O Module Basemodule: ON
```

flashleds

Turns the LEDs for disks in a specific enclosure or bus on or off.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli flashleds` command turns on the LEDs in a specific disk-array enclosure (DPE, iDAE, or DPE) or Fibre Channel bus. The command can help you identify disks in an enclosure or bus.

SYNTAX

```
flashleds -e enclosure-number [-b bus-number] [on|off]
```

OPTIONS

`-e enclosure-number`

Specifies the disk-array enclosure. In each system, if there is a DPE or iDAE, it is enclosure 0. DAEs are numbered sequentially on each bus ascending from 1 (for DPE/iDAE systems) or ascending from 0 (for other system types).

`-b bus-number`

Specifies the bus-number.

`on|off`

Turns the LEDs on or off. If you omit this switch, the command turns the LEDs on.

EXAMPLE

These commands turn on the LEDs for the system that holds SP `ss1_spa` and then turns them off. For that system, it specifies enclosure 1, bus 0.

```
naviseccli -h ss1_sp1 flashleds -e 1 -b 0 on
naviseccli -h ss1_sp1 flashleds -e 1 -b 0 off
```

OUTPUT

There is no output. Errors are printed to `stderr`.

getagent

Gets device names and configurations.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getagent` command returns the names of devices that the destination Unisphere agent controls. The command also displays descriptions of each device with which the Unisphere agent is configured to communicate. If you type the `getagent` command without switches, the CLI displays all values. With switches, the CLI displays only the values.

Each SP is a host (`-h` switch) to which you address the command. The command displays information for the agent running in the SP (the SP agent). Values for irrelevant display fields appear as NA.

Note

If you are running VMware ESX Server and issuing this command to a virtual machine (VM), it cannot target the host agent because the Unisphere Host Agent is supported only on ESX Server.

SYNTAX

```
getagent [-cabinet] [-desc] [-mem] [-model] [-name] [-node]
[-peersig] [-prom] [-rev] [-scsiid] [-serial] [-sig]
[-type] [-ver] [-os] [-spid]
```

OPTIONS

- `-cabinet`
Returns a description of the cabinet type.
 - `-desc`
Returns an ASCII description of the device (taken from the Unisphere agent configuration file).
 - `-mem`
Returns, in MB, the size of the SP's memory. Each SP in the system must have the same amount of memory to make full use of the memory. Total memory can be 8, 16, 32, or 64 MB.
 - `-model`
Returns the model number of the SP.
 - `-name`
Returns the name of the device (taken from the Unisphere agent configuration file).
 - `-node`
Returns the device's location in the file system.
-

Note

Use the returned device name when you specify the device in a command line.

- `-os`
Displays the operating system.
- `-peersig`
Returns the signature of the peer SP.
- `-prom`

Returns the revision of the SP's PROM code. Each SP in the system should be running the same revision of PROM code. PROM code is updated automatically when you update an SP's VNX OE for Block software and choose the reboot option.

-rev

Returns the microcode revision number of the VNX OE for Block software that the SP is running. Each SP in the system will be running the same revision of VNX OE for Block software. You can update an SP's VNX OE for Block software.

-scsiid

Returns the SCSI ID of SP host connection. The SCSI ID number for the SP is determined by switch settings on the SP. For information on these switch settings, see the system installation and maintenance manual.

-serial

Returns the serial number of the system.

-sig

Returns the unique SP signature.

-spid

Returns the SP identifier (A or B).

-type

Returns the model type (deskside or rackmount).

-ver

Returns the version of the agent that is running.

EXAMPLE

This example issues the `getagent` command to system `ss1_spa`.

```
naviseccli -h ss1_spa getagent
```

OUTPUT

The following is a sample output that shows the devices controlled by one agent. Actual output varies depending on the switches you use with the `getagent` command.

Note

For a host that is managing a legacy system, the output may vary depending upon whether or not the host is in a storage group.

Output from an SP agent

```

    Agent Rev: 7.33.0 (4.67)
Name: K10
Desc:
Node: A-FNM00121900314
Physical Node: K10
Signature: 3178501
Peer Signature: 3194985
Revision: 05.33.000.1.467
SCSI Id: 0
Model: VNX8000
Model Type: Rackmount
Prom Rev: 12.50.00
SP Memory: 131072
Serial No: FNM00121900314
SP Identifier: A
Cabinet: SPE9
```

The following output shows the enclosure type for Model VNX5400/VNX5600/VNX5800/VNX7600 Block systems:

```
Cabinet: DPE9
```

The following output shows the enclosure type for a Model VNX8000 Block system:

```
Cabinet: SPE9
```

getall

Gets comprehensive system information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getall` command returns an extensive list of system information, including:

- Agent and host information
- System information
- HBA or switch information
- SP information
- Cache information
- Disk information
- LUN information
- RAID group information
- Storage group information
- SnapView information
- SAN Copy information
- Drive type for enclosures, disks, RAID groups, and LUNs

The `getall` command is a superset of the `getagent` command. If you type the `getall` command without switches, the CLI displays all values. With switches, the CLI displays only the values specified.

Note that the `getall` command and any of its arguments return all information pertinent to the current Unisphere release. The amount of this information may change (usually it increases) with each new Unisphere release. This may affect your custom command scripts that use `getall`.

The command displays information for the agent running in the SP (the SP agent). Values for irrelevant display fields appear as NA.

Some information is not displayed if the feature is not activated or not supported. For example, SAN Copy items are displayed only if SAN Copy software is installed and activated.

SYNTAX

```
getall [-host] [-array] [-hba] [-sp] [-cache] [-disk]
[-drivetype] [-fastcache] [-lun] [-reserved] [-rg] [-sg]
[-snapviews] [-sancopy]
```

OPTIONS

- host
Displays only information about the host and the agent that manage the system. This is a composite of information displayed by the `getagent`, `remoteconfig`, and `ndu` commands.
- array
Displays only information about the system and its hardware components such as fan and link control card (LCC). This is a composite of information displayed by the `getarrayuid`, `arrayname`, `getcrus`, and `initialiazearray` commands.
- hba
Displays information about each switch, HBA and port. It also displays the physical location of the port for a VNX series system. This combines the information that the `alpa`, `port`, `register`, and `getloop` commands display.
- sp
Displays information about this SP only. This is a composite of information displayed by the `getsp`, `getcontrol`, `networkadmin`, and `getsptime` commands.
- cache
Displays information about the cache. This is the same information displayed by the `cache -sp -info` command.
- disk
Displays information about disks and enclosures in the system and new disk spin-down statistics. This is the same information that the `getdisk` command displays.
- drivetype
Displays the drive type of the LUN.
- fastcache
Displays the FAST Cache information.
- lun
Displays only information about each logical unit (host and system) in the system. This is a composite of information displayed by `getsniffer`, `storagegroup`, and `getlun` commands.
- reserved
Displays the reserved LUN pool information.
- rg
Displays only information about RAID groups. This is the same information displayed by `getrg` command.

Note

The `-rg` switch does not display the private RAID group used in thin pools.

- sg
Displays only information about storage groups. This is the same information displayed by the `storagegroup` command.
- snapviews
Displays only information about `snapview`. This is the same information displayed by the `snapview` command.
- sancopy
Displays only information about `sancopy`. This is the same information displayed by the `sancopy` command.

EXAMPLE

This example displays the information about the host and the agent that manage the system.

```
naviseccli getall -host
```

OUTPUT

```

Server IP Address: 10.244.211.40
Agent Rev: 7.33.85 (2.60)
Agent/Host Information
-----
Agent Rev: 7.33.85 (2.60)
Name: K10
Desc:
Node: A-FNM00125200242
Physical Node: K10
Signature: 3309592
Peer Signature: 3313827
Revision: 05.33.000.1.489
SCSI Id: 0
Model: VNX7600
Model Type: Rackmount
Prom Rev: 4.50.00
SP Memory: 65536
Serial No: FNM00125200242
SP Identifier: A
Cabinet: DPE9

Name of the software package: -FASTCache
Revision of the software package: -
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: -ThinProvisioning
Revision of the software package: -
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: -Unisphere
Revision of the software package: -
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: -UnisphereBlock
Revision of the software package: -
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: INTERNAL_USE_ONLY-AutoPilot
Revision of the software package: 05.33.000.1.486
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: INTERNAL_USE_ONLY-FBEAPIX

Revision of the software package: 05.33.000.1.489
Commit Required: NO
Revert Possible: NO

```

```

Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: INTERNAL_USE_ONLY-RALabHosts
Revision of the software package: 05.33.000.1.489
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

Name of the software package: VNX-Block-Operating-Environment
Revision of the software package: 05.33.000.1.489
Commit Required: NO
Revert Possible: NO
Active State: YES
Is installation completed: YES
Is this System Software: NO

```

getarrayuid

Gets the system unique ID.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getarrayuid` command displays the unique ID (UID) of the system that contains the current SP. For a Fibre Channel system this is the system WWN. For an iSCSI system this is also the WWN, but iSCSI also has a target name, the `iqn` identifier, which is a unique identifier. A system with an iSCSI and Fibre Channel combination controller shows both types of IDs.

The SP, HBA, and switch ports also have unique IDs. If you specify one or more SPs, then the CLI displays the unique ID of the systems containing those SPs. Knowing the UID of an SP's system is useful with MirrorView mirroring commands and for other operations.

SYNTAX

```
getarrayuid [-iscsi|-all] [sphostnames]
```

OPTIONS

sphostnames

Are the hostnames of one or more SPs, separated by spaces or commas. If you omit *sphostnames*, then the CLI displays the UID of the SP specified with the `-h` switch.

`-iscsi`

Returns the iSCSI target name (`iqn`) for the system on a combo system.

`-all`

Returns both the system WWN and the iSCSI target name (`iqn`).

Note

If you do not specify any switches, `getarrayuid` returns the system WWN on a system with an iSCSI and Fibre Channel combination controller, as well as on an iSCSI-only system.

EXAMPLE

```
naviseccli -h ssl_spa getarrayuid -all
```

```
Host name      Array UID
Cps42199      50:06:01:60:77:02:C7:A7
1.2.3.4       iqn.1992-04.com.emc.cx.hk192200422
```

OUTPUT

The UID of the system that holds the SP (see above).

getconfig

Returns system type configuration information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getconfig` command displays system type configuration information. If you omit switches, it displays all configuration type information. You can set configuration information with `setconfig`.

Use the `getsp` command to display SP revision and signature information.

SYNTAX

```
getconfig [-ptype] [-que] [-rer] [-tn] [-pg8] [-rep]
```

OPTIONS

`-ptype`

Displays the system package type.

`-que`

Displays the setting on how queue full status is handled. `Yes` means that the VNX OE for Block software will return device status to the operating system as busy, instead of queue full, when the device queue is full.

`-rer`

Displays the number of recovered errors.

`-tn`

Displays the state of the Target Negotiate bit.

`-pg8`

Displays SCSI mode page 8 information.

`-rep`

Displays periodic error report information: the number of reported errors or N/A.

EXAMPLE

```
naviseccli -h ssl_spa getconfig
```

```
Sub-System Package Type: 20
Queue Full Status:      DISABLED
Recovered Errors:       DISABLED
Target Negotiate:       Unknown
Mode Page 8:            DISABLED
Periodic Error Report:  N/A
```


getcontrol

Gets SP performance information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getcontrol` command returns information about the SP performance.

SYNTAX

```
getcontrol [-allib] [-arv] [-busy] [-cbit] [-date] [-day]
[-flush] [-idle] [-read] [-rw] [-sc] [-sf] [-sl] [-time]
[-write] [-all]
```

OPTIONS

`-allib`

Returns the idle and busy ticks for all buses.

`-arv`

Returns the number of times a user request arrived while at least one other request was being processed.

`-busy`

Returns the percentage of time the SP is busy.

`-cbit`

Returns the controller idle and busy ticks.

`-date`

Returns the SP's system date.

`-day`

Returns the SP's system day.

`-flush`

Returns the number of requests to flush the write cache, and the number of write cache blocks flushed.

`-idle`

Returns the percentage of time the SP is idle.

`-read`

Returns the number of reads received by the SP.

`-rw`

Returns the following:

Host read requests

The number of read requests made by the host to the LUN

Host write requests

The number of write requests made by the host to the LUN

Host blocks read

The number of blocks the host read from the LUN

Host blocks written

The number of blocks written from the host to the LUN

`-sc`

Returns the status of the system cache (on/off).

`-sf`

- Returns the status of the system fault LED (on/off).

`-sl`

Returns the status of statistics logging (on/off).

`-time`

Returns the SP's system time.

`-write`

Returns the number of writes received by the SP.

`-all`

In addition to the `getcontrol` command information, it displays the serial number for SP.

EXAMPLE

```
naviseccli -h ss1_spa getcontrol
```

OUTPUT

The following is a sample output. Actual output varies depending upon which switches you use with the `getcontrol` command.

```
System Fault LED: ON
Statistics Logging: ON
SP Read Cache State Enabled
SP Write Cache State Disabled
Max Requests: N/A
Average Requests: N/A
Hard errors: N/A
Total Reads: 0
Total Writes: 0
Prct Busy: 0.72
Prct Idle: 99.2
System Date: 02/06/2013
Day of the week: Wednesday
System Time: 10:02:29
Read_requests: 0
Write_requests: 0
Blocks_read: 0
Blocks_written: 0
Sum_queue_lengths_by_arrivals: 0
Arrivals_to_non_zero_queue: 0
Hw_flush_on: N/A
Idle_flush_on: N/A
Lw_flush_off: N/A
Write_cache_flushes: 70
Write_cache_blocks_flushed: 882
Internal bus 1 busy ticks: N/A
Internal bus 1 idle ticks: N/A
Internal bus 2 busy ticks: N/A
Internal bus 2 idle ticks: N/A
Internal bus 3 busy ticks: N/A
Internal bus 3 idle ticks: N/A
Internal bus 4 busy ticks: N/A
Internal bus 4 idle ticks: N/A
Internal bus 5 busy ticks: N/A
Internal bus 5 idle ticks: N/A
Controller busy ticks: 3651
Controller idle ticks: 505075
```

getcrus

Gets CRU state information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getcrus` command returns state information on all or selected CRUs (customer replaceable units) in the system except for disks (see the `getdisk` command) and selected FRUs/sub FRUs in the system. The `getcrus` command displays the state of the following CRUs: fans, link controller cards (LCCs), storage processors (SPs), battery backup units (referred to as standby power supplies or SPSs), and power supplies (PSs or VSCs). You can use `getcrus` without any switches to get the state information for every CRU (except disks). The `getcrus` command displays the state of management module of FRUs and state of I/O module, CPU module, and DIMM of sub FRUs.

Note

If one of the LCCs is pulled from Bus x , Enclosure x , all data along that bus (after the LCC was pulled) appears as faulted since the agent cannot gather any information along that path. If this occurs, issue the `getcrus` command to the peer SP. The CLI lists the faulted components.

Note

Command switches `-cpua`, `-cpub`, `-dimma`, `-dimmb`, `-ioa`, `-iob`, `-mgmta`, and `-mgmtb` are supported only on Secure CLI.

SYNTAX

```
getcrus [-fana] [-fanb] [-fanc] [-lcca]
[-lccb] [-spa] [-spb]
[-lccrevb] [-lccsnb] [-spsa] [-spsb]
[-lccreva] [-lccsna] [-cablingspsa] [-cablingspsb]
[-encdrivetype] [-speeds] [-cpua] [-cpub] [-ioa] [-iob] [-
dimma]
[-dimmb] [-mgmta] [-mgmtb] [-bma] [-bmb]
[bbua] [bbub] [-all]
```

OPTIONS

- `-bbua`
Returns the status of Battery Backup Unit on SP A.
- `-bbub`
Returns status of Battery Backup Unit on SP B.
- `-bma`
Returns status of Base Module on SP A.
- `-bmb`
Returns status of Base Module on SP B.
- `-cablingspsa`
Returns the cabling status for SPS A (standby power supply A).
- `-cablingspsb`
Returns the cabling status for SPS B (standby power supply B).
- `-cpua`
Returns the state of the CPU module on SP A.
- `-cpub`
Returns the state of the CPU module on SP B.
- `-dimma`
Returns the state of the DIMM module on SP A.
- `-dimmb`
Returns the state of the DIMM module on SP B.
- `-encdrivetype`
Displays the type(s) of drives that are currently allowed to run in the enclosure.

- fana
Returns the state information for fan A.
- fanb
Returns the state information for fan B.
- fanc
Returns the state information for fan C.
- ioa
Returns the state of the I/O modules on SP A.
- iob
Returns the state of the I/O modules on SP B.
- lcca
Returns the state information of the LCC for SP A.
- lccb
Returns the state information of the LCC for SP B.
- lccreva|-lccrevb
Returns the revision of the DPE LCC A, LCC B, or both if you specify both switches.
- lccsna|-lccsnb
Returns the serial number of the DPE LCC A, LCC B, or both if you specify both switches.
- mgmta
Returns the state of management module A.
- mgmtb
Returns the state of management module B.
- spa
Returns the current operational state of SP A. It can be in any of the following states:
 - Present**
The SP is the communication channel you are using to communicate with the chassis.
 - Empty**
Agent cannot talk to the SP because a communication channel specifying the SP is not in the agent's configuration file for the selected host. For example, the SP is connected to a different host than the SP in the communications channel for the chassis.
 - Not Present**
SP that is in the communication channel to the selected chassis has failed or been removed.
 - Removed**
SP was not present when the agent was started.
- spb
Returns the state information for SP B. See -spa description (above).
- spsa
Returns the state information for SPS (standby power supply, backup battery) A.
- spsb
Returns the state information for standby power supply B.
- speeds
Displays the current and maximum speed of enclosures and LCCs within the enclosures. The current speed indicates the speed an enclosure or LCC is currently running at; the maximum speed indicates the highest speed an enclosure or LCC is capable of running at.
- all
Displays the state of all devices.

EXAMPLE

This command gets the information about CRU state.

```
naviseccli -h ss1_spa getcrus
```

OUTPUT

The following example shows the output from the `getcrus -all` command for the VNX5400/VNX5600/VNX5800/VNX7600 Block systems.

```
DPE9 Bus 0 Enclosure 0
Enclosure Drive Type: SAS, NL SAS
Current Speed: 6Gbps
Maximum Speed: 6Gbps
SP A State: Present
SP B State: Present
Bus 0 Enclosure 0 Fan A0 State: Present
Bus 0 Enclosure 0 Fan A1 State: Present
Bus 0 Enclosure 0 Fan B0 State: Present
Bus 0 Enclosure 0 Fan B1 State: Present
Bus 0 Enclosure 0 Power A State: Present
Bus 0 Enclosure 0 Power B State: Present
Bus 0 Enclosure 0 BBU A State: Present
Bus 0 Enclosure 0 BBU B State: Present
Bus 0 Enclosure 0 CPU Module A State: Present
Bus 0 Enclosure 0 CPU Module B State: Present
Bus 0 Enclosure 0 SP A I/O Module 0 State: Present
Bus 0 Enclosure 0 SP A I/O Module 1 State: Empty
Bus 0 Enclosure 0 SP A I/O Module 2 State: Empty
Bus 0 Enclosure 0 SP A I/O Module 3 State: Empty
Bus 0 Enclosure 0 SP A I/O Module 4 State: Empty
Bus 0 Enclosure 0 SP B I/O Module 0 State: Present
Bus 0 Enclosure 0 SP B I/O Module 1 State: Empty
Bus 0 Enclosure 0 SP B I/O Module 2 State: Empty
Bus 0 Enclosure 0 SP B I/O Module 3 State: Empty
Bus 0 Enclosure 0 SP B I/O Module 4 State: Empty
Bus 0 Enclosure 0 DIMM Module A State: Present
Bus 0 Enclosure 0 DIMM Module B State: Present
Bus 0 Enclosure 0 Management Module A State: Present
Bus 0 Enclosure 0 Management Module B State: Present
Bus 0 Enclosure 0 Base Module A State: Present
Bus 0 Enclosure 0 Base Module B State: Present
```

The following example shows the output from the `getcrus -all` command for a VNX8000 Block system.

```
SPE9 Enclosure SPE
Enclosure Drive Type: N/A
Current Speed: N/A
Maximum Speed: N/A
SP A State: Present
SP B State: Present
Enclosure SPE Fan A0 State: Present
Enclosure SPE Fan A1 State: Present
Enclosure SPE Fan A2 State: Present
Enclosure SPE Fan A3 State: Present
Enclosure SPE Fan A4 State: Present
Enclosure SPE Fan B0 State: Present
Enclosure SPE Fan B1 State: Present
Enclosure SPE Fan B2 State: Present
Enclosure SPE Fan B3 State: Present
Enclosure SPE Fan B4 State: Present
Enclosure SPE Power A0 State: Present
Enclosure SPE Power A1 State: Present
Enclosure SPE Power B0 State: Present
Enclosure SPE Power B1 State: Present
Enclosure SPE SPS Module A State: Present
```

```

Enclosure SPE SPS Battery A State: Present
Enclosure SPE SPS Module B State: Present
Enclosure SPE SPS Battery B State: Present
Enclosure SPE SPS Module A Cabling State: Valid
Enclosure SPE SPS Module B Cabling State: Valid
Enclosure SPE CPU Module A State: Present
Enclosure SPE CPU Module B State: Present
Enclosure SPE SP A I/O Module 0 State: Present
Enclosure SPE SP A I/O Module 1 State: Missing
Enclosure SPE SP A I/O Module 2 State: Empty
Enclosure SPE SP A I/O Module 3 State: Empty
Enclosure SPE SP A I/O Module 4 State: Empty
Enclosure SPE SP A I/O Module 5 State: Present
Enclosure SPE SP A I/O Module 6 State: Empty
Enclosure SPE SP A I/O Module 7 State: Empty
Enclosure SPE SP A I/O Module 8 State: Empty
Enclosure SPE SP A I/O Module 9 State: Empty
Enclosure SPE SP A I/O Module 10 State: Empty
Enclosure SPE SP B I/O Module 0 State: Present
Enclosure SPE SP B I/O Module 1 State: Missing
Enclosure SPE SP B I/O Module 2 State: Empty
Enclosure SPE SP B I/O Module 3 State: Empty
Enclosure SPE SP B I/O Module 4 State: Empty
Enclosure SPE SP B I/O Module 5 State: Present
Enclosure SPE SP B I/O Module 6 State: Empty
Enclosure SPE SP B I/O Module 7 State: Empty
Enclosure SPE SP B I/O Module 8 State: Empty
Enclosure SPE SP B I/O Module 9 State: Empty
Enclosure SPE SP B I/O Module 10 State: Empty
Enclosure SPE DIMM Module A State: Present
Enclosure SPE DIMM Module B State: Present
Enclosure SPE Management Module A State: Present
Enclosure SPE Management Module B State: Present
DAE5S Bus 0 Enclosure 0
Enclosure Drive Type: SAS
Current Speed: 6Gbps
Maximum Speed: 6Gbps
Bus 0 Enclosure 0 Power A State: Present
Bus 0 Enclosure 0 Power B State: Present
Bus 0 Enclosure 0 SPS Module A State: Present
Bus 0 Enclosure 0 SPS Battery A State: Present
Bus 0 Enclosure 0 SPS Module B State: Present
Bus 0 Enclosure 0 SPS Battery B State: Present
Bus 0 Enclosure 0 SPS Module A Cabling State: Valid
Bus 0 Enclosure 0 SPS Module B Cabling State: Valid
Bus 0 Enclosure 0 LCC A State: Present
Bus 0 Enclosure 0 LCC B State: Present
Bus 0 Enclosure 0 LCC A Revision: 1.42
Bus 0 Enclosure 0 LCC B Revision: 1.42
Bus 0 Enclosure 0 LCC A Serial #: N/A
Bus 0 Enclosure 0 LCC B Serial #: N/A
Bus 0 Enclosure 0 LCC A Current Speed: 6Gbps
Bus 0 Enclosure 0 LCC B Current Speed: 6Gbps
Bus 0 Enclosure 0 LCC A Maximum Speed: 6Gbps
Bus 0 Enclosure 0 LCC B Maximum Speed: 6Gbps

```

getdisk

Gets disk status.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getdisk` command returns detailed status information about all or selected disks in the system.

A disk module can be in any of the following operational states:

Table 8 Operational states

State	Meaning
Binding	Being bound into a LUN.
Empty	Failed or removed before the agent started running, or the disk wasn't part of a RAID group.
Enabled	Part of a bound LUN.
Expanding	Disk is being added to a RAID group.
Formatting	Being hardware formatted. Generally, modules do not need formatting.
Rebuilding	The data is being rebuilt on a replacement disk module.
Removed	Disk module is removed from the chassis; applies only to a disk module that is part of a LUN.
Low Power	Drive is in low -power state (less than full power); a possible delay can occur with media access.
Reduced Power, Transitioning	Drive is transitioning from a low power/no power to a full power state or vice versa.
Unbound	Ready to be bound into a LUN.

SYNTAX

```
getdisk [-bind] [-capacity] [-actualcapacity] [-userlba]
[-cpct] [-hr] [-hs] [-hw]
[-lun] [-private] [-product] [-rb] [-read] [-rev]
[-serial] [-sectors] [-sig] [-sr] [-state] [-sw]
[-type] [-vendor] [-write] [-rg] [-numluns] [-cpn]
[-rds] [-wrts] [-bytrd] [-bytwrt] [-bndcrs] [-drivetype]
[-tla] [-usercapacity] [-idleticks] [-busyticks] [-speeds]
[-powersavingsdiskcapable] [-powersavingsdiskeligible]
[-powersavingsstate] [-spinstats] [-arrivalswithnonzeroqueue]
[-highsumofseeks] [-idleticks_SPA] [-idleticks_SPB]
[-busyticks_SPA] [-busyticks_SPB] [-que] [-all]
```

OPTIONS

disk-position

Returns status for a specific disk. Refer the *VNX for Block Concepts* guide, section on LUN IDs, unique IDs, and disk IDs to verify the format for specifying disks.

When the *disk-position* is not entered, `getdisk` reports on all disks, regardless of which switches are set. To obtain only the information specified by the optional switches, you must include *disk-position* in the `getdisk` command.

`-actualcapacity`

Returns the actual capacity of the disk.

`-arrivalswithnonzeroqueue`

Returns the cumulative number of requests sent to a disk, while one or more requests were already outstanding.

`-bind`
Returns the percentage of the disk that is bound.

`-bndcrs`
Returns the number of Stripe Boundary Crossings (string value or Unknown).

`-busyticks`
Returns the amount of time that this disk is busy.

`-busyticks_SPA`
Returns the amount of time in 100 ms ticks where the disk was serving requests from SPA, respectively.

`-busyticks_SPB`
Returns the amount of time in 100 ms ticks where the disk was serving requests from SPB, respectively.

`-bytrd`
Returns the number of KB read.

`-bytwr`
Returns the number of KB written.

`-capacity`
Returns the disk capacity in MB.

`-cpn`
Returns the system part number (string value) or N/A.

`-cppct`
Displays the progress of the following three operations:

- Automatic hot spare replacement
- Explicit `copytodisk` command

`-drivetype`
Returns the drive type of the disk.

`-highsumofseeks`
Returns the cumulative seek distance in 512 Byte sectors of all requests to this disk. The seek distance between two consecutive requests is the delta of their corresponding LBAs.

`-hr`
Returns the total number of hard read errors for all the disk modules in the LUN that persisted through all the retries. A growing number of hard errors may mean that one or more of the LUN's disk modules is nearing the end of its useful life.

`-hs`
Returns hot spare status: NO if not a hot spare, the disk position otherwise; see [disk module operational states on page 79](#). The Hot Spare value will be shown as N/A on VNX.

`-hw`
Returns the total number of hard write errors for all the disk modules in the LUN that persisted through all the retries. A growing number of hard errors may mean that one or more of the LUN's disk modules is nearing the end of its useful life.

`-idleticks`
Returns the amount of time that this disk is idle.

`-idleticks_SPA`
Returns the amount of time in 100 ms ticks where the disk had no outstanding requests from SPA, respectively.

`-idleticks_SPB`
Returns the amount of time in 100 ms ticks where the disk had no outstanding requests from SPB, respectively.

`-lun`
Returns the LUN number(s) to which this disk belongs.

`-numluns`

- Returns the number of LUNs bound to this disk. LUNs may span multiple disks.
- powersavingsdiskcapable
 - Displays whether the vendor has approved the drives for spin down. It states YES or NO.
- powersavingsdiskeligible
 - Determines if the drive has met the system requirements (slot location) for transitioning to low-power standby state.
- powersavingsstate
 - Displays the power savings state of the drive.
- private
 - Returns the offset of every partition on the disk.
- que
 - Returns the sum of outstanding requests to the disk. For each newly arriving request, this value gets incremented by the number of currently still outstanding requests to the disk, plus the arriving request itself.
- product
 - Returns the product ID of the disk.
- rb
 - Returns the percentage of the disk that is rebuilt.
- read
 - Returns the total number of read requests made to the disk module. You might find the LUN read information in `getlun` more useful because it is for the entire LUN, and not just for one of the disk modules in the LUN.
- rds
 - Returns the number of read requests.
- rev
 - Returns the product revision number of the disk.
- rg
 - Returns the RAID group (if any) associated with this disk.
- sectors
 - Returns the number of user sectors.
- serial
 - Returns the serial number of the disk.
- sig
 - Returns the bind signature, which was automatically assigned during the bind process. The Bind Signature value will be shown as `N/A, enclosure number, slot number` (for example, `N/A, 0, 1`).
- spinstats
 - Returns the new disk power savings statistics.
- sr
 - Returns the total number of soft read errors for all the disk modules in the LUN that disappeared before all the retries. A growing number of soft errors may indicate that one of the LUN's disk modules is nearing the end of its useful life.
- state
 - Returns the state of the disk.
- sw
 - Returns the total number of soft write errors for all the disk modules in the LUN that disappeared before all the retries. A growing number of soft errors may indicate that one of the LUN's disk modules is nearing the end of its useful life.
- tla
 - Returns the TLA part number.
- type
 - Returns the RAID type of the disk.
- usercapacity
 - Returns the amount of space on the disk that is assigned to bound LUNs.

- userlba
Returns user LBA information.
- vendor
Returns the vendor ID of the disk.
- write
Returns the total number of write requests to the disk module. You might find the LUN write information in `getlun` more useful because it shows the entire LUN, and not just one disk.
- wrts
Returns the number of write requests.
- all
In addition to the `getdisk` command information, it displays maximum queued requests, average queued requests, percent idle, percent busy information, power saving attributes, the percent copied, actual capacity, and LBA of user space.

EXAMPLE

Returns information for disk 1 in disk array enclosure 1 (DAE or DAE2 number 1) on bus 1.

```
naviseccli -h 10.244.211.212 getdisk 0_1_B1
```

OUTPUT

The following is sample output. Actual output varies depending on the switches you use with the `getdisk` command:

```
Bus 0 Enclosure 1 Disk B1
Vendor Id: HITACHI
Product Id: HUC10606 CLAR600
Product Revision: C202
Lun: 16377
Type: N/A
State: Unbound
Hot Spare: N/A
Prct Rebuilt: 16377: 100
Prct Enabled: 16377: 100
Serial Number: PPG9RTPB
Sectors: 1125185536 (549407)
Capacity: 549691
Private: 16377: 65536
Bind Signature: N/A, 0, 1
Hard Read Errors: 0
Hard Write Errors: 0
Soft Read Errors: 0
Soft Write Errors: 0
Read Retries: N/A
Write Retries: N/A
Remapped Sectors: N/A
Number of Reads: 0
Number of Writes: 0
Number of Luns: 1
Raid Group ID: 3
Clariion Part Number: DG118xxxxx3
Request Service Time: N/A
Read Requests: 0
Write Requests: 0
Kbytes Read: 0
Kbytes Written: 0
Stripe Enabledary Crossing: 0
Drive Type: SAS
Clariion TLA Part Number:005000000
User Capacity: 536.530273
Idle Ticks: 0
Busy Ticks: 0
Current Speed: 6Gbps
Maximum Speed: 6Gbps
```

The following is sample output with drive type and state information of disk 0_0_0:

```
naviseccli -h ssl_spa getdisk 0_0_0 -drivetype -state
```

```
Bus 0 Enclosure 0 Disk 0
Drive Type:          SAS
State:               Unbound
```

The following is sample output with the power savings, percent copied, actual capacity, and LBA of user space information:

```
naviseccli -h ssl_spa getdisk 0_4_0 -all
```

```
Bus 0 Enclosure 4 Disk 0
Vendor Id: HITACHI
Product Id: HUC10606 CLAR600
Product Revision: C3B2
Lun: Unbound
Type: N/A
State: Unbound
Hot Spare: N/A
Prct Rebuilt: Unbound
Prct Bound: Unbound
Serial Number: PZG9XKVD
Sectors: 0 (0)
Capacity: 549691
Private: Unbound
Bind Signature: N/A, 4, 0
Hard Read Errors: 0
Hard Write Errors: 0
Soft Read Errors: 0
Soft Write Errors: 0
Read Retries: N/A
Write Retries: N/A
Remapped Sectors: N/A
Number of Reads: 0
Number of Writes: 0
Number of Luns: 0
Raid Group ID: 0
Clariion Part Number: DG118032767
Request Service Time: N/A
Read Requests: 0
Write Requests: 0
Kbytes Read: 0
Kbytes Written: 0
Stripe Boundary Crossing: None
Drive Type: SAS
Clariion TLA Part Number:005049250PWR
User Capacity: 0
Idle Ticks: 0
Busy Ticks: 0
Current Speed: 6Gbps
Maximum Speed: 6Gbps
Queue Max: N/A
Queue Avg: N/A
Prct Idle: Not Available
Prct Busy: Not Available
Hardware Power Savings Qualified: YES
Hardware Power Savings Eligible: YES
Power Savings State: Full Power
Current Power Savings Log Timestamp: N/A
Spinning Ticks: N/A
Standby Ticks: N/A
Number of Spin Ups: N/A
Arrivals with Nonzero Queue: 0
High Sum of Seeks: 0
Idle Ticks SPA: 0
Idle Ticks SPB: 0
```

```

Busy Ticks SPA: 0
Busy Ticks SPB: 0
Queue Length: 0
Prct Copied: N/A
Actual Capacity: 549659
LBA of User Space: 65536

```

getlog

Gets an SP or Event Monitor log.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getlog` command returns the entire SP or Event Monitor log (default), the newest entries to the log, or the oldest entries to the log.

Note

Depending on the size of the log, the `getlog` command may take several minutes to execute.

The command affects only the log of the SP you specify with the `-h` switch in the command.

The `getlog` command requires an active agent.

Note

For a host not attached (through a switch or HBA) to a system, use `getlog -em`.

SYNTAX

```
getlog [-em] [-h] [+n] [-n]
```

OPTIONS

`-em`

Displays the Event Monitor log. If you omit this switch, the CLI displays the Event Monitor log.

`-h`

Displays the `getlog` header.

`+n`

Displays the oldest *n* entries in the log, with the oldest entry first. If there are fewer than *n* entries, the entire log is displayed. *N* must be greater than 0.

`-n`

Displays the newest *n* entries in the log, with the oldest entry first. If there are fewer than *n* entries, the entire log is displayed. *N* must be greater than 0.

If a range of entries is not specified, the entire log is displayed with the oldest entry first.

EXAMPLE

```
naviseccli -h ss1_spa getlog -3
```

```

04/10/1999 12:43:01 SP B(6e5) (FE Fibre loop
operational)[0x00] 0 0
04/10/1999 12:43:19 SP B(6e4) (FE Fibre loop down)
[0x03] 0 0

```

```
04/10/1999 12:43:19 SP B(6e5) (FE Fibre loop
operational) [0x00] 0 0
```

getloop

Gets fibre loop information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getloop` command returns information about the fibre loop.

This command is not valid for IRIX hosts.

Note

This command applies only to Fibre Channel systems.

SYNTAX

```
getloop [-clfo] [-conn] [-dlf] [-icl] [-ple] [-spf]
```

OPTIONS

`-clfo`

Displays whether the command loop has failed over.

`-conn`

Displays the connection number.

`-dlf`

Displays the data loop failover status as follows:

Does system have data loop fail-over: Yes/No

Does system have illegal Cross-Loop: Yes/No

`-icl`

Displays whether there is an illegal cross loop.

`-ple`

Displays whether private loop is enabled.

`-spf`

Displays whether this SP can fail over.

EXAMPLE

```
naviseccli -h ssl_spa getloop
```

```
Has Data Loop Failed Over: NO
Can This SP Failover:      NO
Private Loop Enabled:     YES
Illegal Cross Loop:       NO
Connection #0: Enclosure 0
Connection #1: Enclosure 1
Connection #2: Enclosure 2
```

getlun

Gets LUN information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getlun` command returns information about a LUN and the customer replaceable units (CRUs) that make up that LUN.

Note

If statistics logging is disabled, some fields are not printed (switches marked with * below). Other switches must be used with the `-disk` switch (switches marked with ** below).

If the target LUN is a metaLUN, We recommend that you use the `metalun -list` command (see [metalun -list on page 191](#)) instead of the `getlun` command to return information about the metaLUN. If you use the `getlun` command, only the following fields report data that apply to metaLUNs — **Name, Current owner, Default owner, LUN Capacity (Megabytes), LUN Capacity (Blocks), UID, Is Private, Snapshot List, MirrorView Name if any.**

If you target the LUNs that make up a metaLUN, all fields apply.

If the target LUN is a pool LUN, we recommend that you use the `storagepool -list` command (see [storagepool -list on page 310](#)) instead of the `getlun` command to return information about the pool LUN.

SYNTAX

```
getlun [lun-number] [-aa] [-at] [-bp] [-bread] [-brw]
[-busy] [-bwrite] [-capacity] [-crus] [-default] [-disk]
[-drivetype] [-dsa] [-element] [-ena] [-ff] [-hrdscsi]
[-idle] [-name] [-owner] [-prb]
[-qavg] [-qmax] [-que] [-rc] [-rd3]
[-rchk] [-reads] [-remap] [-rh] [-rhist] [-rr]
[-rwr] [-service] [-sftscsi] [-srcp]
[-stripe] [-totque] [-type] [-upb] [-was]
[-wc] [-wch] [-uid] [-wh] [-whist] [-wr] [-writes]
[-rg] [-state] [-private] [-usage] [-sc] [-addroffset]
[-ismetalun] [-isthinlun] [-ispoolun] [-idleticks] [-
busyticks]
[-lunidleticks] [-luncache] [-lunbusyticks] [-trespass]
[-brwsp] [-rwrsp] [-lunbusytickssp] [-lunidletickssp]
[-nanzq] [-sqlah] [-slst] [-slct] [-etp] [-itp] [-nzcra] [-sor]
[-wcrh] [-fwc] [-fastcache] [issnapmountpoint]
[-idleticks_SPA] [-idleticks_SPB] [-busyticks_SPA] [-
busyticks_SPB]
[-highsumofseeks] [-arrivalswithnonzeroqueue] [-all]
```

OPTIONS*lun-number*

Specifies the logical unit number (LUN) on which to report. If no LUN number is specified, the CLI reports on all LUNs in the system. Refer to *LUN IDs, unique IDs, and disk IDs* in the *VNX for Block Concepts* guide for information about identifying LUN numbers.

`-aa`

Returns the auto-assignment status (enabled/disabled).

`-addroffset`

Displays the starting LBA (logical block address) of the LUN. This address tells you where a LUN begins within a RAID group.

`-at`

Returns auto-trespass status (*enabled* or *disabled*). Shows FAST cache properties.

`-bp*`
Returns the number of blocks prefetched by read cache.

`-bread**`
Returns the number of blocks read for each CRU.

`-brwsp *`
Returns the number of host blocks read and written in SP A and SP B.

`-busy**`
Returns the percentage of time that the disk is busy.

`-busyticks`
Returns the amount of time that the disks on this LUN are busy.

`-bwrite**`
Returns the number of blocks written for each CRU.

`-capacity`
Returns, in MB, the LUN capacity.

`-crus`
Returns the names and states of disks in this LUN.

`-default`
Returns the name of the default owner (SP) of the LUN.

`-disk`
Returns the disk statistics for the disks in the specified *lun-number*. Valid only when statistics logging is enabled.

`-drivetype`
Returns the drive type of the disks in the LUN.

`-dsa`
Displays the status of Dual Simultaneous Access (*dsa*).

`-element`
Returns, in blocks, the stripe size. The stripe size is the number of sectors that the system can read or write to a single disk module in the LUN. The default size is 128 sectors. This size was specified when the LUN was bound. Stripe element size does not apply to a RAID 1 LUN, individual unit, or hot spare.

`-ena**`
Returns the disk status (*enabled/disabled*). Valid only when the `-disk` switch is included in the command.

`-etp*`
Returns the explicit trespasses in SP A and SP B.

`-fastcache`
Displays the FAST Cache properties of the LUN. Enables or disables fast cache usage for each LUN.

`-ff*`
Returns the number of times that a write had to flush a page to make room in the cache.

`-fwc`
Returns the number of write requests that do not have to wait for the allocation of cache pages.

`-hrdscsi`
Returns the total number of hard errors.

`-idle**`
Returns the percentage of time that the disks within the LUN are idle.

`-idleticks`
Returns the amount of time that the disks on this LUN are idle.

`-ismetalun`
Returns information about whether the LUN is a metalUN or not.

`-isthinlun`
Returns information about whether the LUN is a thin LUN or not.

- itp*
Returns the implicit trespasses in SP A and SP B.
- lunbusytickssp*
Returns the amount of time in 10 microsecond increments where any requests to this LUN were outstanding through SPA and SPB, respectively.
- lunidletickssp*
Returns the amount of time in 10 microsecond increments where no requests to this LUN were outstanding through SPA and SPB, respectively.
- luncache
Returns LUN offline (cache dirty) condition information for the specified LUN. See also the `luncache` command.
- mirrorname
Returns the MirrorView mirror name.
- name
Returns the LUN name. This switch applies only to systems running a version of VNX OE for Block software that supports storage group commands. LUN names are displayed in the UI. You cannot use the LUN name as a CLI command argument.

Note

The `-name` switch applies only to shared systems.

- nanzq*
Returns the number on arrivals with nonzero queue.
- nzcra*
Returns the nonzero request count arrivals in SP A and SP B.
- owner
Returns the name of the SP that currently owns the LUN. If neither SP owns the LUN a value of Unknown is returned.
- prb
Returns the percentage of the LUN that is rebuilt.
- private
Returns information on all private LUNs in the system.
- que
Returns the queue length, a number.
- rc
Returns the LUN's read cache state (`alwaysenabled`).

Note

The read SP cache is enabled by default while binding LUNs on flash drives. It is enabled for all RAID types.

- reads**
Returns the number of reads received for each CRU.
- readtime
Returns the cumulative read time in microseconds.
- rg
Returns the RAID group ID.
- rhist*
Returns Read Histogram information. Read Histogram is a system of 10 locations that contain the number of reads. Element n of the system contains the number of reads that were larger than or equal to 2^{n-1} and less than 2^n blocks in size. For example, the first bucket contains the number of 1 block reads, the second contains the number of 2-3 block reads, the third contains the number of 4-7 block reads, and

so on. `-rhist` also displays read histogram overflows, which are the number of I/O operations that were larger than 512 blocks.

`-rwrsp*`
Returns the number of host read and write requests in SP A and SP B.

`-sc`
Returns the LUN capacity as stripe count.

`-sftscsi`
Returns the total number of soft (correctable) errors.

`-slct*`
Returns the statistics logging current time.

`-slst*`
Returns the statistics logging start time.

`-snapshot`
Returns the number of SnapView snapshots (copy images on which a SnapView session is active).

`-sor*`
Returns the sum of outstanding requests in SP A and SP B.

`-sqlah*`
Returns the sum of queue lengths on arrivals high.

`-srcp`
Returns the read cache configuration.

`-state`
Returns the state of the LUN. Valid states are: Expanding, Defragmenting, Faulted, Transitional, or bound.

`-stripe`
Returns the number of times an I/O crossed a stripe boundary on a RAID 6, RAID 5, RAID 0, or RAID 1/0 LUN.

`-totque`
Returns the total queue length, a number.

`-type`
Returns the RAID type of the LUN as follows:

```
RAID 0 - nonredundant individual access array
RAID 1 - mirrored pair
RAID 3 - single-disk parity
RAID 5 - distributed parity
RAID 6 - double distributed parity
RAID 1/0 - mirrored RAID 0 group
Disk - individual unit
Hot Spare - hot spare
```

`-uid`
Returns the LUN unique ID.

`-upb*`
Returns the number of prefetched blocks not used by the read cache.

`-usage`
Returns the usage for a private LUN. For standard LUNs the usage displays as Unknown.

`-was`
Returns the minimum block size request that bypasses cache.

`-wc`
Returns the LUN's write cache state (enabled or disabled).

Note

The write SP cache is enabled by default while binding LUNs on flash drives. It is enabled for all RAID types.

- wch*
Returns the number of write requests that do not have to wait for the allocation of cache pages.
- wcrh
Returns the number of write requests to a previously written address location whose data is still resident in cache. Write rehits form a subset of write cache hits since they overwrite already existing data in the write cache.
- wh*
Returns the write hit information, if statistics logging is enabled.
- whist*
Returns write histogram information. Write Histogram is an system of 10 locations that contain the number of writes. Element n of the system contains the number of writes that were larger than or equal to 2^{n-1} and less than 2^n blocks in size. For example, the first bucket contains the number of 1 block writes, the second contains the number of 2-3 block writes, the third contains the number of 4-7 block writes, and so on. `whist` also displays the write histogram overflows, which are the number of writes that were larger than 512 blocks.
- writes**
Returns the number of writes received for each CRU.
- writetime
Returns the cumulative write time in microseconds.
- all
Lists the `getlun` command information along with the FAST Cache information.

EXAMPLE #1

This example retrieves the LBA (logical block address), LUN capacity, and alignment offset information.

```
naviseccli -h ssl_spa getlun -addroffset -offset -capacity
```

```
LOGICAL UNIT NUMBER 0
Offset:                N/A
LUN Capacity(Megabytes): 1024
LUN Capacity(Blocks):  2097152
Address Offset:        0

LOGICAL UNIT NUMBER 1
Offset:                N/A
LUN Capacity(Megabytes): 5120
LUN Capacity(Blocks):  10485760
Address Offset:        2097152

LOGICAL UNIT NUMBER 2
Offset:                N/A
LUN Capacity(Megabytes): 10240
LUN Capacity(Blocks):  20971520
Address Offset:        12582912
```

EXAMPLE #2

This example retrieves the cumulative read time and write time.

```
naviseccli -h ssl_spa getlun 0 -readtime -writetime
```

```
Cumulative Read Time (microseconds): 87937
Cumulative Write Time (microseconds): 0
```

EXAMPLE #3

This example retrieves information about LUN number 0: (See sample listing following)

```
naviseccli -h ssl_spa getlun 0
```

OUTPUT

The following is a sample output. Actual output varies depending on the switches you use with the getlun command.

```
LOGICAL UNIT NUMBER 0
Prefetch size (blocks) = N/A
Prefetch multiplier = N/A
Segment size (blocks) = N/A
Segment multiplier = N/A
Maximum prefetch (blocks) = N/A
Prefetch Disable Size (blocks) = N/A
Prefetch idle count = N/A

Prefetching: N/A Prefetched data retained N/A

Read cache configured according to specified parameters.

Total Hard Errors: 0
Total Soft Errors: 0
Total Queue Length: 0
Name LUN 0
Minimum latency reads N/A

RAID Type: RAID5
RAIDGroup ID: 0
State: Bound
Stripe Crossing: 0
Element Size: 128
Current owner: SP B
Offset: N/A
Auto-trespass: DISABLED
Auto-assign: DISABLED
Write cache: ENABLED
Read cache: ENABLED
Idle Threshold: N/A
Idle Delay Time: N/A
Write Aside Size: 0
Default Owner: SP B
Rebuild Priority: N/A
Verify Priority: N/A
Prct Reads Forced Flushed: N/A
Prct Writes Forced Flushed: N/A
Prct Rebuilt: 100
Prct Bound: N/A
LUN Capacity(Megabytes): 51200
LUN Capacity(Blocks): 104857600
UID: 60:06:01:60:E8:90:32:00:EF:14:8D:66:13:4E:E2:11
Bus 0 Enclosure 4 Disk 19 Queue Length: 0
Bus 0 Enclosure 4 Disk 0 Queue Length: 0
Bus 0 Enclosure 4 Disk 2 Queue Length: 0
Bus 0 Enclosure 4 Disk 4 Queue Length: 0
Bus 0 Enclosure 4 Disk 6 Queue Length: 0
Bus 0 Enclosure 4 Disk 7 Queue Length: 0
Bus 0 Enclosure 4 Disk 19 Hard Read Errors: 0
Bus 0 Enclosure 4 Disk 0 Hard Read Errors: 0
Bus 0 Enclosure 4 Disk 2 Hard Read Errors: 0
```

```

Bus 0 Enclosure 4 Disk 4 Hard Read Errors: 0
Bus 0 Enclosure 4 Disk 6 Hard Read Errors: 0
Bus 0 Enclosure 4 Disk 7 Hard Read Errors: 0
Bus 0 Enclosure 4 Disk 19 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 0 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 2 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 4 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 6 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 7 Hard Write Errors: 0
Bus 0 Enclosure 4 Disk 19 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 0 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 2 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 4 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 6 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 7 Soft Read Errors: 0
Bus 0 Enclosure 4 Disk 19 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 0 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 2 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 4 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 6 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 7 Soft Write Errors: 0
Bus 0 Enclosure 4 Disk 19 Enabled Bus 0 Enclosure 4 Disk 0 Enabled
Bus 0 Enclosure 4 Disk 2 Enabled Bus 0 Enclosure 4 Disk 4 Enabled
Bus 0 Enclosure 4 Disk 6 Enabled Bus 0 Enclosure 4 Disk 7 Enabled
Is Private: NO
Snapshots List: Not Available
MirrorView Name if any: Not Available

```

getresume

Displays resume information for system devices.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getresume` command displays the properties (resume information) for system devices, such as the enclosure display board, link control card (LCC), power supply (PS), storage processor (SP), SAN personality card, central processing unit module (CPU), I/O module, management module, and the standby power supply (SPS). You can display resume information for a specific device or for all devices.

The product serial number, product part number, and product revision are displayed only for the SPE enclosure on VNX systems. This information is displayed only when you use the `-all` switch. Only Secure CLI supports use of this switch.

SYNTAX

```

getresume [-sp] [-lcc busNumber enclosureNumber [lcca|lccb]]
          [-ps busNumber enclosureNumber|xpe [psa|psb]]
          [-pc] [-mp] [-sps] [-vendorinfo] [-all]

```

OPTIONS

`-all`

Displays resume information for all devices.

`-lcc busNumber enclosureNumber [lcca|lccb]`

Displays resume information for the link control card. If you omit `lcca` or `lccb`, information is returned for link control cards A and B.

`-mp`

Displays resume information for the enclosures.

`-pc`

- Displays resume information for the SAN personality card.
- `-ps busNumber enclosureNumber| xpe [psa|psb]`
Displays resume information for the power supply. If you omit psa or psb, information is returned for power supplies A and B.
- `-sp`
Displays resume information for SPs.
- `-sps`
Displays resume information for the standby power supply (SPS) that is associated with the SP you specify in the command line.

EXAMPLE

This example displays information stored in the resume PROM for power supplies A and B (bus 0, enclosure 0).

```
naviseccli -h ssl_spa getresume -ps 0 0
```

```
Bus 0 Enclosure 0
  Power A
    EMC Part Number: N/A
    EMC Artwork Revision: N/A
    EMC Assembly Revision: N/A
    EMC Serial Number: N/A
    Vendor Part Number: SG9006-710G
    Vendor Artwork Number: 010
    Vendor Assembly Number: A10
    Vendor Serial Number: AC7111210221
    Vendor Name: ACBEL POLYTECH INC.
    Location of Manufacture: TANG XIA TOWN, DONG GUAN, CHINA

    Year of Manufacture: 2011
    Month of Manufacture: 03
    Day of Manufacture: 30
    Assembly Name: Dual +12V P/S & Cooling Module.

    Programmable Name: 1sa :2sa :3sa :4sa

    Programmable Revision: 5.21:3.31:3.21:2.04

Bus 0 Enclosure 0
  Power B
    EMC Part Number: N/A
    EMC Artwork Revision: N/A
    EMC Assembly Revision: N/A
    EMC Serial Number: N/A
    Vendor Part Number: SG9006-710G
    Vendor Artwork Number: 010
    Vendor Assembly Number: A10
    Vendor Serial Number: AC7111210334
    Vendor Name: ACBEL POLYTECH INC.
    Location of Manufacture: TANG XIA TOWN, DONG GUAN, CHINA

    Year of Manufacture: 2011
    Month of Manufacture: 03
    Day of Manufacture: 30
    Assembly Name: Dual +12V P/S & Cooling Module.

    Programmable Name: 1sa :2sa :3sa :4sa

    Programmable Revision: 5.33:3.31:3.21:2.04
```

OUTPUT

The following example shows sample output of resume information for I/O carrier device 0 on both SPs.

```

Storage Processor A
I/O Carrier 0
EMC Part Number          204-012-901D
EMC Artwork Revision     N/A
EMC Assembly Revision    D04
EMC Serial Number        CF2JY063400197
Vendor Part Number       N/A
Vendor Artwork Number    N/A
Vendor Assembly Number   N/A
Vendor Serial Number     N/A
Vendor Name              CELESTICA
Location of Manufacture  THAILAND
Year of Manufacture      2006
Month of Manufacture     9
Day of Manufacture       15
Assembly Name            IO Carrier Device
Programmable Name        NVRAM0:FLASH0:NVRAM1
Programmable Revision    0.01:2.07:0.01

Storage Processor B
I/O Carrier              0
EMC Part Number          204-012-901D
EMC Artwork Revision     N/A
EMC Assembly Revision    D04
EMC Serial Number        CF2JY063400174
Vendor Part Number       N/A
Vendor Artwork Number    N/A
Vendor Assembly Number   N/A
Vendor Serial Number     N/A
Vendor Name              CELESTICA
Location of Manufacture  THAILAND
Year of Manufacture      2006
Month of Manufacture     9
Day of Manufacture       14
Assembly Name            IO Carrier Device
Programmable Name        NVRAM0:FLASH0:NVRAM1
Programmable Revision    0.01:2.07:0.01
Enclosure SPE

Chassis/Midplane
EMC Part Number: 100-562-266
EMC Artwork Revision: N/A
EMC Assembly Revision: A01
EMC Serial Number: HK100072600036
Vendor Part Number: N/A
Vendor Artwork Number: N/A
Vendor Assembly Number: N/A
Vendor Serial Number: N/A
Vendor Name: N/A
Location of Manufacture: Hopk,MA USA
Year of Manufacture: 2007
Month of Manufacture: 07
Day of Manufacture: 11
Assembly Name: DREADNOUGHT DVT CHASSIS
Programmable Name: Micro 11.33
Programmable Revision: Micro 11.33
EMC Product Serial Number: XXXXXXXXXXXXXXXX
EMC Product Part Number: 900-XXX-XXX
EMC Product Revision: XX.XX.XX
Bus 0 Enclosure 0

Chassis/Midplane
EMC Part Number: 100-561-622
EMC Artwork Revision: D05

```

```

EMC Assembly Revision: A08
EMC Serial Number: FCNST064401287
Vendor Part Number: N/A
Vendor Artwork Number: N/A
Vendor Assembly Number: N/A
Vendor Serial Number: N/A
Vendor Name: FOXCONN, SHENZHEN, CHINA
Location of Manufacture: LONGHUA TOWN, SHENZHEN, CHINA
Year of Manufacture: 2006
Month of Manufacture: 10
Day of Manufacture: 30
Assembly Name: 4GB FC DAE STILETTO W/FLANGES
Programmable Name: Micro 11.33
Programmable Revision: Micro 11.33
Bus 1 Enclosure 0

Chassis/Midplane
EMC Part Number: 100-561-622
EMC Artwork Revision: D05
EMC Assembly Revision: A08
EMC Serial Number: FCJST064500188
Vendor Part Number: N/A
Vendor Artwork Number: N/A
Vendor Assembly Number: N/A
Vendor Serial Number: N/A
Vendor Name: FOXCONN, SHENZHEN, CHINA
Location of Manufacture: LONGHUA TOWN, SHENZHEN, CHINA
Year of Manufacture: 2006
Month of Manufacture: 11
Day of Manufacture: 09
Assembly Name: 4GB FC DAE STILETTO W/FLANGES
Programmable Name: Micro 11.33
Programmable Revision: Micro 11.33

```

getrg

Gets RAID group information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getrg` command returns information about the specified RAID group.

If no RAID group is specified, the command returns information about all RAID groups.

SYNTAX

```

getrg rg [-type] [-state] [-disks] [-drivetype] [-lunlist]
[-maxd] [-maxl] [-tcap] [-ucap] [-lusc] [-pod] [-prcntdf]
[-prcntex] [-exdisks] [-lunex] [-legal] [-powersavingssetting]
[-powersavingseligible] [-isrginstandbystate] [-elementsize]
[-all]

```

OPTIONS

rg

Is the RAID group identification number. *rg* specifies a number between 0 and the maximum number of RAID groups supported by the system.

`-disks`

Returns the state of disks in the RAID group.

`-drivetype`

Returns information about what drive type the RAID group consists of.

- exdisks
Returns information about which disks are expanding. Displays N/A for not expanding.
- legal
Returns information about specific RAID type LUNs that you can bind for a specific RAID group. How you can bind the new LUNs depends on the number of LUNs that already exist and upon the number of disks in the RAID group.
- lunex
Returns information about which LUNs are expanding. Displays NO for not expanding.
- lunlist
Returns a list of LUNs that are currently in the RAID group.
- lusc
Returns a free contiguous group of unbound segments in blocks.
- maxd
Returns the maximum number of disks that are allowed in the RAID group.
- maxl
Returns the maximum number of LUNs that are allowed in the RAID group.
- pod
Returns the priority of defragmentation/expansion operations. The Defrag/Expand priority value is always shown as N/A on VNX.
- powersavingssetting
Sets the power savings setting ON or OFF. The default value for a nonconfigured RAID group is OFF.
- powersavingseligible
Determines the eligibility of a RAID group for the power savings option. A RAID group is eligible for power savings when all the disks in the RAID group are eligible for power savings.
- isrginstandbystate
Returns the power savings state of the RAID group when all the disks participating in RAID group are in power savings mode.
- prcntdf
Returns the percent of defragmentation that is complete. The Percent defragmented value is always shown as N/A on VNX.
- prcntex
Returns the percent of expansion that is complete. The Percent expanded value is always shown as N/A on VNX.
- state
Returns the RAID group state, which may be one of the following: Invalid, Valid_luns, Halted, and Busy.
- tcap
Returns the raw and logical capacity of the RAID group in blocks.
- type
Returns the RAID group type.
- ucap
Returns the free (unbounded) capacity of the RAID group in blocks.
- all
Lists the RAID group information.
- elementsize
Displays the element size of a RAID group. The element size of a LUN is always equal to the element size of the RAID Group that holds the LUN.

EXAMPLE

Retrieves information about RAID group 1.

```
naviseccli -h ss1_spa getrg 1 -all
```


OUTPUT

The following is a sample output. The actual output varies depending on the switches you use with the `getrg` command.

```
RaidGroup ID: 1
RaidGroup Type: r5
RaidGroup State: Valid_luns
List of disks: Bus 0 Enclosure 1 Disk 24
Bus 0 Enclosure 1 Disk 23
Bus 0 Enclosure 1 Disk 22
Bus 0 Enclosure 1 Disk 21
Bus 0 Enclosure 1 Disk 20
List of luns: 0 1 2 3
Max Number of disks: 16
Max Number of luns: 256
Raw Capacity (Blocks): 2799769600
Logical Capacity (Blocks): 2239815680
Free Capacity (Blocks,non-contiguous): 2223038464
Free contiguous group of unbound segments: 2223038464
Defrag/Expand priority: N/A
Percent defragmented: N/A
Percent expanded: N/A
Disk expanding onto: N/A
Lun Expansion enabled: NO
Legal RAID types: r5
Power Savings Setting: OFF
RAID GROUP Power Savings Eligible: NO
Is RAID GROUP in Power Savings Mode: NO
Drive Type: SAS
Element Size: 1024
```

If you request the `-elementsiz` option explicitly, the following output is generated:

```
RaidGroup ID: 1
Element Size: 128
```

The following is a sample output of the `getrg -drivetype` command.

```
RaidGroup ID: 1
Drive Type: SATA Flash

RaidGroup ID: 5
Drive Type: NL SAS
```

getsniffer

Gets background verify reports.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getsniffer` command retrieves background verify reports. You can retrieve reports for a specific LUN, all LUNs in a RAID group, or all LUNs in a system.

Note

The `getsniffer` command does not support thin LUNs.

SYNTAX

```
getsniffer lun-number|-rg rg-number]| [-all] [-curr]
[-rec] [-alltot]
```

OPTIONS

-all

Retrieves reports on all LUNs in the system.

Note

When you specify the `-all` switch, it may generate a sizeable amount of data and take a considerable amount of time to generate the data.

lun-number

Specifies the logical unit on which to report. *lun-number* specifies the logical unit number.

-rg *rg-number*

Retrieves reports on all LUNs in the RAID group.

-alltot

Retrieves and displays the report of historical totals of all full unit verifies.

-curr

Retrieves and displays the report of the currently running full unit verify.

-rec

Retrieves and displays the report of the most recently completed full unit verify.

EXAMPLE

Retrieves the report of the active full unit verify operation for LUN 0.

```
naviseccli -h ssl_spa getsniffer 0
```

OUTPUT

The following is a sample output. Actual output varies depending on the switches you use with the `getsniffer` command.

```
VERIFY RESULTS FOR UNIT 0
Currently Running Full Unit Verify
-----
                Corrected                Uncorrectable
Checksum errors
0
    0
Write Stamp errors
0
    0
Time Stamp errors
0
    0
Shed Stamp errors
0
    0
Coherency errors
0
    0
Multi Bit CRC errors
0
    0
LBA Stamp errors
0
    0
Media errors
```

```

0
  0
Soft Media errors                0

Most Recently Completed Full Unit Verify
-----
                Corrected                Uncorrectable
Checksum errors
0
  0
Write Stamp errors
0
  0
Time Stamp errors
0
  0
Shed Stamp errors
0
  0
Coherency errors
0
  0
Multi Bit CRC errors
0
  0
LBA Stamp errors
0
  0
Media errors
0
  0
Soft Media errors                0

Historical Total of All Full Unit Verifies (3 passes)
-----
                Corrected                Uncorrectable
Checksum errors
0
  0
Write Stamp errors
0
  0
Time Stamp errors
0
  0
Shed Stamp errors
0
  0
Coherency errors
0
  0
Multi Bit CRC errors
0
  0
LBA Stamp errors
0
  0
Media errors
0
  0
Soft Media errors                0

```

getsp

Gets the SP revision and serial number information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

Displays SP identification information. If you omit switches, it displays all SP configuration information.

Note

Use the `getsp` command to display SP revision and signature information.

SYNTAX

```
getsp [-type] [-sig] [-psig] [-rev] [-ser] [-mem] [-id]
```

OPTIONS

`-type`

Displays the SP type.

`-sig`

Displays the SP unique signature.

`-psig`

Displays the SP peer unique signature.

`-rev`

Displays the SP revision number.

`-ser`

Displays the SP serial number.

`-mem`

Displays the SP memory size.

`-id`

Displays the SP SCSI ID if available, otherwise displays N/A.

EXAMPLE #1

```
naviseccli -h ssl_spa getsp
```

```
SP A
Cabinet: SPE9
Signature For The SP: 3178501
Signature For The Peer SP: 3194985
Revision Number For The SP: 05.33.000.1.467
Serial Number For The SP: FCNMS120700013
Memory Size For The SP: 131072
SP SCSI ID if Available: 0

SP B
Cabinet: SPE9
Signature For The SP: 3194985
Signature For The Peer SP: 3178501
Revision Number For The SP: 05.33.000.1.467
Serial Number For The SP: FCNMS121200069
Memory Size For The SP: 131072
SP SCSI ID if Available: 0
```

EXAMPLE #2

The following command displays the enclosure type for VNX5400/VNX5600/VNX5800/VNX7600 Block systems.

```
naviseccli -h ss1_spa getsp -type
```

```
SP A
Cabinet: DPE9

SP B
Cabinet: DPE9
```

EXAMPLE #3

The following command displays the enclosure type for VNX8000 Block systems.

```
naviseccli -h ss1_spa getsp -type
```

```
SP A
Cabinet:           SPE9

SP B
Cabinet:           SPE9
```

getsptime

Gets the date and time setting on each SP system clock.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli` command displays the time on the system clocks in one or both storage processors. The VNX OE for Block software ensures that the system clocks in both SPs are close but necessarily identical.

SYNTAX

```
getsptime [-spa|-spb]
```

OPTIONS

`-spa` or `-spb`

Specifies from which SP to display the date and time.

EXAMPLE

```
naviseccli -h 10.35.50.11 -user xxxxxx -password xxxxxx -scope 0
getsptime
```

```
Time on SP A: 08/25/10 10:09:28
Time on SP B: 08/25/10 10:09:40
```

OUTPUT

The SP date and time.

getspuptime

Gets the length of time since the SPs were last rebooted.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

This `naviseccli` command determines the length of time since the Storage Processor (SPA or SPB) on the VNX system was last rebooted.

SYNTAX

```
getspuptime [-spa|-spb]
```

OPTIONS

`-spa` or `-spb`

Specifies the SP for which to display the time since the last reboot. Not specifying an option displays the uptime for both SPs.

EXAMPLE

```
naviseccli -h a.b.c.d -user xxxxxx -password xxxxxx -scope 0
getspuptime
```

OUTPUT

```
SP A Uptime: 51 days 10 hours 15 minutes
SP B Uptime: 51 days 9 hours 45 minutes
```

“N/A” is shown if the peer information could not be retrieved.

getunusedluns

Returns the number of LUNs not used by any storage group or by any replication applications.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli getunusedluns` command displays all the unused LUNs, which are those the storage group and replication applications are not using.

SYNTAX

```
getunusedluns
```

EXAMPLE

```
naviseccli -h ss1_spa getunusedluns
```

OUTPUT

```
RaidGroup ID: 0
List of luns: 20

RaidGroup ID: 1
List of luns: 4094 4095 4092 4093 4090 21

RaidGroup ID: 2
List of luns: 10 11 12 13 14

RaidGroup ID: 3
List of luns: 4091 16 17 18 19

RaidGroup ID: 4
List of luns: 0 1 2 3 4
```

hotsparepolicy -list

Lists the hot spare policies for all classes of disks on the array.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hotsparepolicy -list` command lists the hot spare policies for all classes of disks on the array. Disk class is defined by disk type-capacity combination.

SYNTAX

```
hotsparepolicy -list [-allentries]
```

OPTIONS

`-allentries`

Displays all entries, including those for any disk classes that ever appeared on the array.

EXAMPLE

This example lists all the hot spare policies for the classes of disks currently residing on the array.

```
naviseccli hotsparepolicy -list
```

OUTPUT

The following is a sample output. Actual output varies depending upon which switch you use with the command.

```
Policy ID:                1
Disk Type:                SAS
Raw Capacity (MBs):      274845
Number of Disks:         30
Unused disks for hot spares: 15
Recommended Ratio of Keep Unused  1/30
Ratio of Keep Unused     1/30
Number to Keep Unused:   1
Recommended:             Yes

Policy ID:                2
Disk Type:                NL SAS
Raw Capacity (MBs):      2748450
Number of Disks:         5
Unused disks for hot spares: 0
Recommended Ratio of Keep Unused  1/30
Ratio of Keep Unused     1/30
Number to Keep Unused:   1
Recommended:             Yes
```

hotsparepolicy -set

Sets the hot spare policy.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hotsparepolicy -set` command sets the hot spare policy.

SYNTAX

```
hotsparepolicy -set id -keep1unusedper ratio [-o]
```

OPTIONS*id*

Specifies the ID of the hot spare policy.

-keep1unusedper *ratio*

Specifies hot spare policy ratio. The allowable value for *ratio* is an integer equal to or larger than 0. For example, 0 means no hot spare, 1 means all disks are kept as hot spare, 15 means to keep one disk per 15 disks are kept as hot spares.

-o

Executes the command without prompting for confirmation.

EXAMPLE

This example changes the hot spare policy ratio to 1/15.

```
naviseccli hotsparepolicy -set 1 -keep1unusedper 15
You are about to change the hotspare policy, do you want to change
the ratio to 1/15 for SAS drives with 274845 MBS capacity for hot
spare use? [Y/N] ? (Y)
```

OUTPUT

None if the command succeeds; status or error information if it fails.

inserttestevent

Inserts an event into the event monitor log to let you verify the accuracy of an event monitor template.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli inserttestevent` command inserts a dummy event of code in the event log. The event is code `0x2003` and its description is `Test Event - internal use only`. The event is in the local event queue for processing.

To verify that the monitoring agent will respond as specified (for example, issue email, send a page) use the `responsetest` command.

SYNTAX

```
inserttestevent
```

EXAMPLE

For `ss1_spa`, this command writes a test event into the log file.

```
naviseccli -h ss1_spa inserttestevent
```

ioportconfig -config -maxvlans

Sets the maximum number of VLANs per port.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli ioportconfig -config -maxvlans` command sets the maximum number of VLANs per port.

SYNTAX

```
ioportconfig -config [-maxvlans 8|16] [-delay delay] [-o]
```

OPTIONS

`-maxvlans 8 | 16`

Sets the maximum number of VLANs per port to either 8 or 16.

`-delay delay`

Delays the coordinated reboot between SPs for the specified value. The value ranges from 0 to 1140 seconds. The default value is 360 seconds.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This example shows how to set the maximum number of VLANs per port to 16.

```
naviseccli -h ssl_spa -user user11 -password pswd11 -scope 0
ioportconfig -config -maxvlans 16
```

OUTPUT

```
WARNING: This operation will change the maximum VLAN ports. The
operation will reboot the SPs.
DO YOU WISH TO CONTINUE (Y/N)?
```

ioportconfig -list

Displays information about all I/O modules and corresponding I/O ports.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli ioportconfig -list` command displays the information about all I/O modules and I/O ports on VNX series systems.

SYNTAX

```
ioportconfig -list -iomodule "onboard"|"basemodule"|slot_num
[-sp a|b] [label] [-pportid] [-portrole] [-portusage]
[-porttype] [-portstate] [-portsubstate] [-isportpersisted]
[-lportid] [-all]
```

OPTIONS

`-list`

Displays information about all I/O modules and I/O ports (for both SPs).

`-iomodule "onboard"|"basemodule"|slot_num`

Displays information about the module and all I/O ports that exist on the given module. If the `slot_num` is not provided, the system returns an error.

Note

`"basemodule"` is not supported on VNX8000 Block systems.

`-sp a | b`

- Specifies the SP, either A or B. If an `-sp` switch is not specified, then information about the I/O module with the given `slot_num` for both SPs will be shown.
- `-label`
Displays the IO Module Label information.
- `-pportid`
Displays the physical port number.
- `-portrole`
Displays whether the port is operating as a front-end or back-end port. If not persisted, this field displays an error `Uninitialized`.
- `-portusage`
Indicates how the port is being used. If the port is a MirrorView port, it displays `Special`. If the port is initialized, and is not the MirrorView port, it displays `Normal`. If the port is not initialized, the state is `Uninitialized`. Refer to *EMC MirrorView/Asynchronous Command Line Interface (CLI) Reference* for information about the MirrorView ports for the various systems.
- `-porttype`
Displays the port transport type (Fibre Channel or iSCSI).
- `-portstate`
Displays the state of the port. If not persisted, this field displays `Uninitialized`.
- `-portsubstate`
Provides a detailed description of why the port is in a particular state.
- `-isportpersisted`
Indicates whether the port is currently persisted.
- `-lportid`
Shows the logical port ID.

EXAMPLE

This example shows information about all I/O modules and corresponding I/O ports.

```
naviseccli -h ssl_spa -user user11 -password pswd11 -scope 0
ioportconfig -list -all
```

OUTPUT

```
Information about each I/O module(s) on SPA:

SP ID:                A
I/O Module Slot:     0
I/O Module Type:     Unknown
I/O Module State:    Empty
I/O Module Substate: Not Present
I/O Module Power state: Not Available
I/O Carrier:        No
I/O Module Label:    Unknown

SP ID:                A
I/O Module Slot:     1
I/O Module Type:     Fibre Channel
I/O Module State:    Present
I/O Module Substate: Good
I/O Module Power state: On
I/O Carrier:        No
I/O Module Label:    8 Gb Fibre

Information about each port on this I/O module:
Physical Port ID:    0
Port Role:          FE
Logical Port ID:    0
Port Usage:         Special
Port Type:          Fibre Channel
Port State:         Enabled
```

```

Port Substate:                Good
Is Persisted:                 Yes

Physical Port ID:             1
Port Role:                    FE
Logical Port ID:              1
Port Usage:                   Normal
Port Type:                    Fibre Channel
Port State:                   Enabled
Port Substate:                Good
Is Persisted:                 Yes

Physical Port ID:             2
Port Role:                    FE
Logical Port ID:              2
Port Usage:                   Normal
Port Type:                    Fibre Channel
Port State:                   Enabled
Port Substate:                Good
Is Persisted:                 Yes

Physical Port ID:             3
Port Role:                    FE
Logical Port ID:              3
Port Usage:                   Normal
Port Type:                    Fibre Channel
Port State:                   Enabled
Port Substate:                Good
Is Persisted:                 Yes

SP ID:                        A
I/O Module Slot:              2
I/O Module Type:              iSCSI
I/O Module State:             Present
I/O Module Substate:          Good
I/O Module Power state:       On
I/O Carrier:                  No
I/O Module Label:             10 GbE v3

Information about each port on this I/O module:
Physical Port ID:             0
Port Role:                    FE
Logical Port ID:              4
Port Usage:                   Special
Port Type:                    iSCSI
Port State:                   Enabled
Port Substate:                Good
Is Persisted:                 Yes

Physical Port ID:             1
Port Role:                    FE
Logical Port ID:              5
Port Usage:                   Normal
Port Type:                    iSCSI
Port State:                   Enabled
Port Substate:                Good
Is Persisted:                 Yes

SP ID:                        A
I/O Module Slot:              3
I/O Module Type:              iSCSI
I/O Module State:             Present
I/O Module Substate:          Good
I/O Module Power state:       On
I/O Carrier:                  No
I/O Module Label:             10 GbE BaseT

Information about each port on this I/O module:

```

```

Physical Port ID:      0
Port Role:            FE
Logical Port ID:      6
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

Physical Port ID:      1
Port Role:            FE
Logical Port ID:      7
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

SP ID:                A
I/O Module Slot:      4
I/O Module Type:      iSCSI
I/O Module State:     Present
I/O Module Substate:  Good
I/O Module Power state: On
I/O Carrier:          No
I/O Module Label:     1 GbE iSCSI/TOE

Information about each port on this I/O module:
Physical Port ID:      0
Port Role:            FE
Logical Port ID:      8
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

Physical Port ID:      1
Port Role:            FE
Logical Port ID:      9
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

Physical Port ID:      2
Port Role:            FE
Logical Port ID:      10
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

Physical Port ID:      3
Port Role:            FE
Logical Port ID:      11
Port Usage:           Normal
Port Type:            iSCSI
Port State:           Enabled
Port Substate:        Good
Is Persisted:         Yes

SP ID:                A
I/O Module Slot:      Base Module
I/O Module Type:      SAS
I/O Module State:     Present

```

```

I/O Module Substate:      Good
I/O Module Power state:   On
I/O Carrier:              No
I/O Module Label:         N/A

Information about each port on this I/O module:
Physical Port ID:         0
Port Role:                 BE
Logical Port ID:          0
Port Usage:                Normal
Port Type:                 SAS
Port State:                Enabled
Port Substate:            Good
Is Persisted:             Yes

Physical Port ID:         1
Port Role:                 BE
Logical Port ID:          1
Port Usage:                Normal
Port Type:                 SAS
Port State:                Missing
Port Substate:            SFP is missing
Is Persisted:             Yes

Information about each I/O module(s) on SPB:

SP ID:                     B
I/O Module Slot:           0
I/O Module Type:           Unknown
I/O Module State:          Empty
I/O Module Substate:       Not Present
I/O Module Power state:    Not Available
I/O Carrier:               No
I/O Module Label:          Unknown

SP ID:                     B
I/O Module Slot:           1
I/O Module Type:           Fibre Channel
I/O Module State:          Present
I/O Module Substate:       Good
I/O Module Power state:    On
I/O Carrier:               No
I/O Module Label:          8 Gb Fibre

Information about each port on this I/O module:
Physical Port ID:         0
Port Role:                 FE
Logical Port ID:          0
Port Usage:                Special
Port Type:                 Fibre Channel
Port State:                Enabled
Port Substate:            Good
Is Persisted:             Yes

Physical Port ID:         1
Port Role:                 FE
Logical Port ID:          1
Port Usage:                Normal
Port Type:                 Fibre Channel
Port State:                Enabled
Port Substate:            Good
Is Persisted:             Yes

Physical Port ID:         2
Port Role:                 FE
Logical Port ID:          2
Port Usage:                Normal
Port Type:                 Fibre Channel

```

```

Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 3
Port Role: FE
Logical Port ID: 3
Port Usage: Normal
Port Type: Fibre Channel
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

SP ID: B
I/O Module Slot: 2
I/O Module Type: iSCSI
I/O Module State: Present
I/O Module Substate: Good
I/O Module Power state: On
I/O Carrier: No
I/O Module Label: 10 GbE v3

Information about each port on this I/O module:
Physical Port ID: 0
Port Role: FE
Logical Port ID: 4
Port Usage: Special
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 1
Port Role: FE
Logical Port ID: 5
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

SP ID: B
I/O Module Slot: 3
I/O Module Type: iSCSI
I/O Module State: Present
I/O Module Substate: Good
I/O Module Power state: On
I/O Carrier: No
I/O Module Label: 10 GbE BaseT

Information about each port on this I/O module:
Physical Port ID: 0
Port Role: FE
Logical Port ID: 6
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 1
Port Role: FE
Logical Port ID: 7
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

```

```

SP ID: B
I/O Module Slot: 4
I/O Module Type: iSCSI
I/O Module State: Present
I/O Module Substate: Good
I/O Module Power state: On
I/O Carrier: No
I/O Module Label: 1 GbE iSCSI/TOE

Information about each port on this I/O module:
Physical Port ID: 0
Port Role: FE
Logical Port ID: 8
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 1
Port Role: FE
Logical Port ID: 9
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 2
Port Role: FE
Logical Port ID: 10
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 3
Port Role: FE
Logical Port ID: 11
Port Usage: Normal
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

SP ID: B
I/O Module Slot: Base Module
I/O Module Type: SAS
I/O Module State: Present
I/O Module Substate: Good
I/O Module Power state: On
I/O Carrier: No
I/O Module Label: N/A

Information about each port on this I/O module:
Physical Port ID: 0
Port Role: BE
Logical Port ID: 0
Port Usage: Normal
Port Type: SAS
Port State: Enabled
Port Substate: Good
Is Persisted: Yes

Physical Port ID: 1
Port Role: BE

```

```

Logical Port ID:          1
Port Usage:              Normal
Port Type:               SAS
Port State:              Missing
Port Substate:           SFP is missing
Is Persisted:            Yes

Information about maximum VLANs per port:
Max VLANs per port:      8

```

ioportconfig -list -maxvlans

Displays information about the maximum number of VLANs per I/O port.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `ioportconfig -list -maxvlans` command displays the information about VLANs per I/O port on VNX series systems.

SYNTAX

```
ioportconfig -list [-maxvlans]
```

OPTIONS

-maxvlans

Shows the maximum number of VLANs per port.

EXAMPLE

This example shows information for VLANs.

```
naviseccli -h ssl_spa -user user11 -password pswd11 -scope 0
ioportconfig -list -maxvlans
```

OUTPUT

```

Information about maximum VLANs per port:
Max VLANs per port:  8

```

ioportconfig -persist

Writes the port configuration information to persistent memory.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli ioportconfig -persist` command writes all port configuration information on the subsystem to persistent memory and reboots both SPs. You can specify the delay value for writing the port configuration information. If you do not specify the delay, then the system takes the default value of 360 seconds.

SYNTAX

```
ioportconfig -persist [-delay delay] [-o]
```

OPTIONS

-delay *delay*

Delays the coordinated reboot between SPs for the specified value. The value ranges from 0 to 1140 seconds. The default value is 360 seconds.

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -h ioportconfig -persist -delay 20
```

OUTPUT

```
This operation will cause both SPs to reboot in a coordinated
fashion. DO YOU WISH TO CONTINUE (Y/N)?
```

luncache -clear

Clears the LUN offline (cache dirty) condition.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You can verify the offline condition of a LUN, using the `luncache -list` command.

DESCRIPTION

The `naviseccli luncache` command with the `-clear` function clears the LUN offline (cache dirty) condition. You can clear the LUN offline condition for a specific LUN, or for all LUNS on the system.

When a storage processor (SP) fails, writes destined for a LUN that have not yet been flushed to disk place the LUN in an offline (cache dirty) condition. LUNs in an offline condition are inaccessible until pending writes for the LUN are cleared.

Note

When you execute `luncache -clear`, all pending writes on the LUN are lost. You must issue the command to the SP that owns the LUN. The `luncache -clear` command does not support thin LUNs.

SYNTAX

```
luncache lun-number -clear [-o]
```

OPTIONS

lun-number

Specifies the LUN for which to clear the LUN offline condition. If you do not specify a LUN number, clears the LUN offline condition for all LUNS on the system.

-o

Executes the command without prompting for confirmation.

EXAMPLE

This command clears the LUN offline condition for the specified LUN.

```
naviseccli -h ss1_spa luncache 9 -clear
```

```
Bringing the LUN online will cause all pending writes
on the LUN to be lost.
```

```
Do you want to bring the LUN online now? (y/n)
```

OUTPUT

None if the command succeeds; status or error information if it fails.

luncache -list

Returns LUN offline (cache dirty) condition information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli luncache` command with the `-list` function returns the LUN offline (cache dirty) condition information. You can return information for a specific LUN or for all LUNS on the system.

When a storage processor (SP) fails, writes destined for a LUN that have not yet been flushed to disk place the LUN in an offline (cache dirty) condition. LUNs in an offline condition are inaccessible until pending writes for the LUN are cleared.

Note

You must issue the command to the SP that is the default owner of the LUN which is not necessarily the current owner. You can clear the LUN offline condition using `luncache -clear`. The `luncache -list` command does not support thin LUNs.

SYNTAX

```
luncache lun-number -list
```

OPTIONS

lun-number

Indicates the LUN for which to display LUN offline condition information. If you do not specify a LUN number, LUN offline condition information is returned for all LUNS on the system.

EXAMPLE

This command displays LUN offline condition information for the specified LUN.

```
naviseccli -h ss1_spa luncache 9 -list
```

OUTPUT

```
LUN Offline (Cache Dirty Condition):    YES
```

managedby

Displays the application that manages the system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

You use the `managedby` command for purposes of the VSS Provider. The VSS Provider is used in conjunction with the SnapView functionality, to provide backup capabilities.

Note

The command returns an output value of `Navi Manager (Unisphere)` or `None`. A value of `None` is synonymous to `Unisphere`.

SYNTAX

```
managedby
```

EXAMPLE

The system is managed by Unisphere.

```
naviseccli -h ss1_spa managedby
```

OUTPUT

```
Managed By: Navi Manager
```

managefiles -delete

Deletes SP logging files on the specified system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli managefiles` command with the `-delete` function lets you delete SP logging files from supported directories.

Note

The dumps and logs directories are supported for `managefiles`.

You can specify the files you want to delete by using the `-all` or `-file` switches, or you can omit switches and display a list of files, and then choose a file from the list.

SYNTAX

```
managefiles -delete [-all] [-file filenames] [-o]
```

OPTIONS

`-all`

Deletes all the files from the supported directory.

`-file filenames`

Deletes the specified files from the supported directory.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command shows you all file index numbers and names. The files appear in the format shown below. You then enter the file you want to delete, by either its index number or name.

```
naviseccli -h ss1_spa managefiles -delete
```

Index	Size	Last modified	Filename
0	0	10/26/2000 17:39:53	naviagent_Oct-26-00_13-36-17.log
1	138	10/27/2000 12:37:19	naviagent_Oct-26-00_13-48-40.log
2	0	10/27/2000 13:20:53	naviagent_Oct-27-00_08-48-38.log
3	0	10/27/2000 13:29:57	naviagent_Oct-27-00_09-29-57.log
4	0	10/27/2000 18:31:42	naviagent_Oct-27-00_11-06-09.log
5	0	10/27/2000 18:41:37	naviagent_Oct-27-00_14-41-37.log

```
Enter files to be deleted with index separated by comma
(1,2,3,4-5) OR a range (1-3) OR enter 'all' to delete all
file OR 'quit' to quit> 2
```

```
Files selected to be deleted are
naviagent_Oct-27-00_08-48-38.log
Do you want to continue(y/n)[n]? (y/n)? y
File naviagent_Oct-27-00_08-48-38.log is deleted.
```

OUTPUT

None if the command succeeds; status or error information if it fails.

managefiles -list

Lists SP logging files in supported directories.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli managefiles` command with the `-list` function lets you display a list of SP logging files in supported directories.

Note

The dumps and logs directories are supported for `managefiles`.

SYNTAX

```
managefiles -list
```

EXAMPLE

This command shows you a list of files in the supported directories, for the system you specify.

```
naviseccli -h ssl_spa managefiles -list
```

OUTPUT

Index	Size	Last modified	Filename
0	0	10/26/2000 17:39:53	naviagent_Oct-26-00_13-36-17.log
1	138	10/27/2000 12:37:19	naviagent_Oct-26-00_13-48-40.log
2	0	10/27/2000 13:20:53	naviagent_Oct-27-00_08-48-38.log
3	0	10/27/2000 13:29:57	naviagent_Oct-27-00_09-29-57.log
4	0	10/27/2000 18:31:42	naviagent_Oct-27-00_11-06-09.log
5	0	10/27/2000 18:41:37	naviagent_Oct-27-00_14-41-37.log

managefiles -retrieve

Retrieves SP logging files from a remote system to the local system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli managefiles` command with the `-retrieve` function lets you retrieve SP logging files from supported directories on a remote system to the local system.

Note

The dumps and logs directories are supported for `managefiles`.

You can specify the files you want to retrieve by using the `-all` or `-file` switches, or you can omit switches and display a list of files, and then choose a file from the list.

SYNTAX

```
managefiles -retrieve [-path path] [-all] [-file filenames]
[-o]
```

OPTIONS

`-path path`

Specifies the path on the local system where you want to copy the file. If you do not use this switch, the file is copied to the current directory.

`-all`

Retrieves all files from the supported directory.

`-file filenames`

Retrieves the specified files from the supported directory.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command shows you all file index numbers and names. The files appear in the format shown below. You then enter the file you want to retrieve, by either its index number or name.

```
naviseccli -h ssl_spa managefiles -retrieve
```

Index	Size	Last modified	Filename
0	0	10/26/2000 17:39:53	naviagent_Oct-26-00_13-36-17.log
1	138	10/27/2000 12:37:19	naviagent_Oct-26-00_13-48-40.log
2	0	10/27/2000 13:20:53	naviagent_Oct-27-00_08-48-38.log
3	0	10/27/2000 13:29:57	naviagent_Oct-27-00_09-29-57.log
4	0	10/27/2000 18:31:42	naviagent_Oct-27-00_11-06-09.log
5	0	10/27/2000 18:41:37	naviagent_Oct-27-00_14-41-37.log

```
Enter files to be retrieved with index seperated by comma
(1,2,3,4-5) OR a range (1-3) OR enter 'all' to retrieve
all file OR 'quit' to quit> 2
```

```
Files selected to be retrieved are
naviagent_Oct-27-00_08-48-38.log
```

```
Do you want to continue(y/n)[n]? (y/n)? y
```

```
File naviagent_Oct-27-00_08-48-38.log is retrieved to the current
directory.
```

OUTPUT

None if the command succeeds; status or error information if it fails.

ndu -list

Describes installed SP driver software packages.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must be able to log in to the host running VNX CLI for Block.

DESCRIPTION

The `ndu` command provides nondisruptive upgrades on VNX series systems. This command lets you install and upgrade system software without disrupting any I/O from attached servers. Third-party software installation is not supported. This means you can use `ndu` to install VNX OE for Block software bundles or enablers.

Use the `-list` function with no switches to display all the information about all software packages.

The `naviseccli ndu` command with the `-list` function and no switches displays information about all installed software packages. The format looks like this:

```
Name of the software package:      name
Revision of the software package: n.nn
Commit Required:                  yes or no or
                                   already committed
Revert Possible:                  yes or no
Active State:                     yes or no
Is installation complete           yes or no or not
                                   applicable.
Is this System Software:          yes or no
```

For information on one package, use the `-name` switch. For a subset of package information, use one or more additional switches.

After listing installed driver packages, you can perform other `ndu` (non-disruptive upgrade) steps, such as `ndu -commit` or `ndu -revert`.

SYNTAX

```
ndu -list [-name [name]] [-rev] [-iscommitable] [-isrevertable]
[-isinactive] [-iscomplete] [-issystem] [-gen]
```

OPTIONS

`-name [name]`

Without the *name* argument, displays only the names of packages; with *name*, displays information on the package *name*.

`-gen`

Displays the generation of the package.

`-rev`

Displays the package revision.

`-iscommitable`

Displays *yes* if the package can be committed, *no* if it cannot.

`-isrevertable`

Displays *yes* if the package can be reverted (it is not committed), *no* if it cannot be.

`-isinactive`

Displays *yes* if the package is active, *no* if it is not.

`-iscomplete`

Displays *yes* if the package file contains a complete set of files, *no* or not applicable if it does not.

`-issystem`

Displays *yes* if the package is a factory-supplied system driver, *no* if it is not.

EXAMPLE

For the SP with hostname `payroll_storage_SPA`, this command lists all installed driver packages. For sample output, see previous page.

```
naviseccli -h payroll_storage_SPA ndu -list
```

OUTPUT

See previous page. If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

ndu -install

Transfers SP software driver packages to the system private LUN (PSM LUN).

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must have a user entry in the Unisphere agent configuration file.

Preinstallation validation checks

Preinstallation validation checks identify unsupported or unsafe installation conditions. You initiate the validation checks functionality when you issue the `ndu -install` command. The validation checks run in the background, prior to installing the software. If a validation check fails, the CLI displays the error and terminates the installation. You can choose to display all validation checks as the functionality executes by specifying the `-verbose` switch, otherwise the CLI only displays failures that prevent installation. The following responses exist for each validation check:

Table 9 Validation check responses

Response	Definition
Successful	Validation check is successful. Installation proceeds.
Warning	Conditions exist that may need correction. Installation proceeds.
Failure	Conditions exist that require correction before installation can proceed. Installation terminates.

Note

The preinstallation validation checks identify a series of unsupported or unsafe installation conditions, but cannot guarantee a successful install.

In addition to the user access requirements for the `ndu` command, the validation check has associated configuration requirements that you must follow:

- Create the required user accounts on the system (refer to the *VNX for Block Concepts* guide, section Getting Started with Secure CLI) . You can create the user accounts using the Secure CLI syntax.
- Create a security file on the host. If you do not create a security file, the CLI prompts you for a valid username, password and scope:

```
Security File does not exist
Please enter security information to proceed:
  Enter User (Existing user on system)
  Enter Password
  Enter Scope (0 - global [default]; 1 - local
```

Note

If the system is uninitialized, user credential information is not required. Therefore, a security file is not required and the CLI does not prompt you for a username, password and scope.

You can also issue validation checks without installing or upgrading software (see `ndu -runrules`).

Note

Use the `naviseccli ndu -list` command to display installed package information. If the `ndu -install` command succeeds, the SPs restart.

DESCRIPTION

The `ndu` command with the `-install` function transfers files from media to the system, queries and displays information about the packages, and then installs or upgrades selected software packages.

The `naviseccli ndu` command `-install` function transfers one or more SP driver packages from a user-accessible file system to the system private storage LUN (PSM). Media should be present before you issue this command.

You should use a Windows-based management software (either Unisphere or VNX CLI for Block) to install software.

Before starting a non-disruptive software installation, record the read and write cache sizes because they will be set to zero.

Before the SP starts a nondisruptive software installation, it disables the caches and sets their sizes to zero. If the write cache is full and I/O is heavy, disabling the cache may take over an hour because the cached data must be written to disk. After the data is written, the installation starts.

When you install new SP software using the CLI, the only way to determine when the installation is finished is to issue periodic `ndu -status` commands until the CLI shows the operation is completed.

When the installation is complete, restore the cache sizes to their original sizes if possible. You may not be able to use the original sizes because the new software requires more memory than the version that it replaced.

The software prompts for information as needed; then it installs or upgrades the specified software packages and restarts the SPs. The SPs then load and run the new packages. After successful installation, it deletes the files from the system.

You can install more than one package with one `ndu` command.

Note

When you install an upgrade (that is, a newer version of an installed package), you must install all the software packages you want to use in the same command. For example, if you are upgrading SnapView in a system that has SnapView, Access Logix, and VNX OE for Block software installed, then you must upgrade all three using one `ndu -install` command. When you install a new package of the same revision as other existing packages, you may install only that package and not the others.

The CLI takes the following actions in order:

- If it can find the packages, it transfers them without interaction.

- Unless you used `-force` (which prevents the prompt and is equivalent to the `all` response), the software examines each package and displays the following information for confirmation:

```

Item number:                n
Name of the software package: driver-name
Revision of the software package: n.nn
Already Installed Revision: NO or YES
Installable                 YES or NO

Enter Item number of the packages to be installed separated
by blanks. Enter 'all' for all packages and
'quit' to quit without installing:

```

Specify the package(s) you want installed. If you make an invalid choice, the CLI does not install any package.

Note

Do not change the configuration; for example, binding new LUNs or expanding metaLUNs while a software install is in progress. Such operations will be rejected; delay them until after the software upgrade is complete and committed.

SYNTAX

```
ndu -install pathlist [-delay seconds] [-force] [-gen]
[-verbose]
```

OPTIONS

pathlist

Specifies filenames, with full pathnames, of the software packages to be installed. You can use a Windows or UNIX pathname (for example, `C:\temp` or `/usr/bin`). Enclose the pathname list in quotes and separate multiple names by a space. If you are installing a newer version of an existing package, you must install all other packages (as explained above).

`-delay delay`

Specifies the delay time that allows attached hosts enough time after SP B has finished rebooting to rescan and mark paths available to SP B, before SP A reboots. The minimum allowed delay is zero seconds and the maximum allowed is 1140 seconds. In most cases, the EMC recommended NDU delay of 360 seconds should be used. However, some host configurations (such as those using Veritas DMP) may require a longer delay.

`-force`

Installs without user interaction. If any package is not installed for any reason, the software displays an error message and continues installing other packages specified in the command.

`-gen`

Displays the generation of the package.

`-verbose`

Displays all results of the preinstallation validation checks. Without this switch, only errors that prevent installation (failures) are displayed.

EXAMPLE

```
naviseccli -h ssl_spa ndu -install
```

```
VNX-Block-Operating-Environment-05.33.000.1.539.pbu
```

```
Item number: 0
Name of the software package: VNX-Block-Operating-Environment
Revision of the software package: 05.33.000.3.539
Already Installed Revision 05.33.000.3.486
Installable YES
Disruptive upgrade: NO

The requested package(s) will be installed. Do you wish to proceed?
: (y/n)?y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Not supported error message is printed to stderr. Other error message(s):

```
NDU_NOT_SUPPORTED
NDU_CANNOT_FIND_FILE
NDU_CANNOT_XFER_FILE
INVALID_PACKAGE
```

ndu -runrules

Executes the preinstallation validation checks.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must have a user entry in the Unisphere agent configuration file.

DESCRIPTION

The `naviseccli ndu` command with the `-runrules` function executes the preinstallation validation checks, without performing software installation (see `ndu -install`).

Note

The `ndu -runrules` command has associated configuration requirements for the preinstallation validation checks functionality. See Preinstallation validation checks under [ndu -install on page 119](#).

Note

Use the `naviseccli ndu -install` command to both run the preinstallation validation checks and install software.

SYNTAX

```
ndu -runrules [pathlist] [-listrules] [-verbose]
```

OPTIONS

pathlist

Specifies filenames, with full pathnames, of the software packages to be committed.

`-listrules`

Displays a list of preinstallation validation checks that would run for the specified host. The validation checks do not execute. To run the validation checks, do not include this switch.

`-verbose`

Displays all results of the preinstallation validation checks. Without this switch, only errors that prevent installation (failures) are displayed.

EXAMPLE

For the SP with hostname `payroll_storage_SPA`, this command executes the preinstallation validation checks.

```
naviseccli -h ssl_spa ndu -runrules
```

OUTPUT

If the preinstallation validation checks detect failures, the failures are displayed.

ndu -commit

Commits (makes permanent) an installed system driver package.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must have a user entry in the Unisphere agent configuration file.

DESCRIPTION

The `naviseccli ndu` command with the `-commit` function commits an installed software package. Every package does not require commit.

If a package needs to be committed and it is not committed, all new features of the newly installed package may not be available. A committed package cannot be reverted. It can be uninstalled.

Note

You cannot create any new RAID groups or bind any new LUNs until you commit VNX OE for Block and Access Logix software.

Use the `naviseccli ndu -listcommand` to display installed package information.

SYNTAX

```
ndu -commit namelist
```

OPTIONS

namelist

Lists package names of the software packages to be committed.

EXAMPLE

For the SP with hostname `payroll_storage_SPA`, this command commits the installed package `mypackage`.

```
naviseccli -h ssl_spa ndu -commit mypackage
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

ndu -status

Returns the status of the ndu install, commit, or revert operations.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must have a user entry in the Unisphere agent configuration file.

DESCRIPTION

The `naviseccli ndu` command with the `-status` function reports the progress of the last `ndu -install`, `-commit`, or `-revert` function executed.

Note

Use the `naviseccli ndu -list` command to display installed package information.

SYNTAX

```
ndu -status [-clear]
```

OPTIONS

`-clear`

Clears the status of the last executed `ndu` command.

EXAMPLE

For the SP with hostname `payroll_storage_SPA`, this command reports the status of the `ndu -install`, `-commit`, or `-revert` function.

```
naviseccli -h ssl_spa ndu -status
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

ndu -revert

Returns to (restores functionality of) the previous revision of an installed system driver.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

You must have a user entry in the Unisphere agent configuration file.

DESCRIPTION

The `naviseccli ndu` command with the `-revert` function restores the functionality of the previous revision of an installed driver package. A committed package cannot be reverted, nor can a package that had no previous revision installed. After the command succeeds, the SPs will restart.

If you omit the `-o` (override) switch, the CLI prompts for confirmation:

```
Revert operation will revert package-name from both SPs. Do you
still want to revert. (y/n)?
```

Answer **y** to revert; answer **n** to cancel the command.

Note

Use the `naviseccli ndu -list` command to display installed package information.

SYNTAX

```
ndu -revert [-delay seconds] namelist [-o]
```

OPTIONS

`-delay seconds`

Specifies the delay time that allows attached hosts enough time after SP B has finished rebooting to rescan and mark paths available to SP B, before SP A reboots. The minimum allowed delay is zero seconds and the maximum allowed is 1140 seconds. In most cases, the EMC recommended NDU delay of 360 seconds should be used. However, some host configurations (such as those using Veritas DMP) may require a longer delay.

`namelist`

Lists package names of the software packages to be reverted.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For the SP with hostname `payroll_storage_SPA`, this command reverts to the previous revision of `mypackage` for both SPs.

```
naviseccli -h ssl_spa ndu -revert mypackage
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
NDU_CANNOT_FIND_FILE
NDU_CANNOT_XFER_FILE
INVALID_PACKAGE
```

networkadmin -get

Lists network name and address information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli networkadmin` command with `-get` function without switches lists all the network information for an SP. This command supports IPv6 network communication to the management ports of systems with VNX OE for Block version 04.28 or later. If you do not specify the IP address (either IPv4 or IPv6), then by default IPv4 information is displayed. The default IP address is IPv4. If the `-ipv4` and the `-ipv6` switches are specified, the information is displayed for both IP addresses with the IPv4 information listed first.

Issue the command to the SP for which this information is needed.

Issuing this command with one or more switches displays the information based on the specified switch. The line `Storage Processor:` is always displayed. This command displays information in the following format:

```
Storage Processor:           SP-id (A or B)
Storage Processor Network Name: SP-network-name
Storage Processor IP Address:  SP-IP-address
Storage Processor Subnet Mask: SP-subnet-mask
Storage Processor Gateway Address: SP's-gateway-system address
```

Note

If you need to change an SP network name or address, then use the `networkadmin` command with the `-set` function described in [networkadmin -set on page 131](#).

SYNTAX

```
networkadmin -get [-sp a|b] [-name] [-portid portid [-vportid
vportid]] [-vlanid]
[-ipv4 [-mode] [-address] [-subnetmask] [-gateway]]
[-ipv6 [-mode] [-address] [-globalprefix] [-gateway]]
[-speed] [-auto] [-capablespeeds] [-requested-speed]
[-linkstatus] [-all]
```

OPTIONS

`-sp a|b`

Displays the properties of the specified SP. The default is the properties of the connected SP.

`-portid portid`

Displays the management port ID. If you do not specify the port ID, the system displays the information of all management ports.

`-vportid vportid`

Displays the virtual port ID associated with the management port. If you do not specify the virtual port ID, the system displays the information of all virtual ports.

`-vlanid`

Displays the virtual LAN (VLAN) ID of the management port/virtual port. If the VLANs are disabled, the system displays `Disabled`.

`-name`

Displays the SP's network name.

`-ipv4`

Displays all the IPv4 information. This is the default IP address.

`-ipv6`

Displays all the IPv6 information.

`-address`

Displays the SP's network IP address (IPv4/IPv6 address).

`-gateway`

Displays the IP address (IPv4/IPv6 address) of the SP's gateway system.

`-subnetmask`

Displays the SP's IPv4 subnet mask.

`-globalprefix`

Displays the SP's IPv6 global prefix.

`-mode`

Displays the IPv4/IPv6 status. The statuses are automatic, manual, and disabled.

`-speed`

Displays the current port speed for a specified SP.

-auto (CX4 series only)

Indicates whether the auto-negotiate feature is turned on or off. The CLI returns a `yes` value if the feature is turned on, a `no` value if the feature is turned off, and a `not supported` value if the feature is not supported.

Note

If you have the auto-negotiate feature turned on, the speed value that is generated in the output will not be `auto`. The value reflects the speed at which the port is currently operating. If you have the Auto-Negotiate feature turned on and you specify the `-requestedspeed` switch, the Requested Value displays a value of `auto`.

-capablespeeds (CX4 series only)

Returns a list of valid speed values for each port.

-requestedspeed (CX4 series only)

Displays the requested speed value.

-linkstaus (CX4 series only)

Displays the link status of the management port.

-all (CX4 series only)

Displays all information for the network admin command.

Note

The command switches `-speed`, `-auto`, `-capablespeeds`, `-requestedspeed`, `-linkstatus`, `-portid`, `-vportid`, `-vlanid`, `-ipv4`, `-ipv6`, and `-all` are supported only with Secure CLI.

EXAMPLE #1

```
naviseccli -h ss1_spa networkadmin -get -speed
```

```
Storage Processor SP A
Management Port Settings
Link Status: Link-up
Current Speed: 100 Mbps/half duplex
```

EXAMPLE #2

```
naviseccli -h ss1_spa networkadmin -get -capablespeeds -auto
```

```
Storage Processor : SP A
Management Port Settings:
Link Status:      Link-up
Current Speed:    100 Mbps/half duplex
Auto-Negotiate:   No
Capable Speeds:  10 Mbps half/full duplex
                  100 Mbps half/full duplex
                  1000 Mbps half/full duplex Auto
```

EXAMPLE #3**Note**

The IPv6 information contains IPv6 in the header to distinguish the output from the IPv4 information.

```
naviseccli -address 10.5.2.12 -user a -password a -scope 0
networkadmin -get -ipv4 -ipv6
```

```
Storage Processor: SP A
Storage Processor Network Name: lumpy-spa
Storage Processor IP Address: 10.5.2.12
Storage Processor Subnet Mask: 255.255.255.0
Storage Processor Gateway Address: 10.5.2.1
Storage Processor IPv6 Mode: Automatic
Storage Processor IPv6 Address:
3ffe:80c0:22c:4d:20d:56ff:fec5:bff4
Storage Processor IPv6 Global Prefix: 3ffe:80c0:22c:4d::
Storage Processor IPv6 Local Address:
fe80::20d:56ff:fec5:bff4
Storage Processor IPv6 Gateway Address:
fe80::20a:8bff:fe5a:967c
```

EXAMPLE #4

```
naviseccli -h ssl_spa networkadmin -get -all
```

```
Storage Processor: SP A
Storage Processor Network Name: kirk-spa

Port ID: 0
Management Port Settings:
Link Status: Link-Up
Current Speed: 100Mbps/full duplex
Requested Speed: Auto
Auto-Negotiate: YES
Capable Speeds: 10Mbps half/full duplex
                 100Mbps half/full duplex
                 1000Mbps half/full duplex
                 Auto

Virtual Port ID: 0
VLAN ID: 102
Storage Processor IP Mode: Manual
Storage Processor IP Address: 10.5.2.178
Storage Processor Subnet Mask: 255.255.255.0
Storage Processor Gateway Address: 10.5.2.1
Storage Processor IPv6 Mode: Automatic
Storage Processor IPv6 Address:
3ffe:80c0:22c:47:260:1600:3ce0:151c
Storage Processor IPv6 Global Prefix:
3ffe:80c0:22c:47::
Storage Processor IPv6 Link-local Address:
fe80::260:1600:3ce0:151c
Storage Processor IPv6 Gateway Address:
fe80::20a:8bff:fe5a:967c
```

OUTPUT

See the previous page. If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

If the command switches `-speed`, `-auto`, `-capabilities`, `-requested-speed`, `-link-status`, and `-all` are not supported on the system then an `Invalid command line parameters` error message is printed to `stderr`.

If IPv6 does not support the system an error is displayed: `IPv6 is not supported for this platform`.

IPv6 is enabled but a value for one of the IPv6 values has not been set so, the value will be displayed as an empty string.

networkadmin -mib

Configures the SNMP services on the system SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli networkadmin` command with the `-mib` function enables or disables the processing of SNMP MIB read requests on an SP and sets the SNMP community string. This command with no option lists the current status of the SNMP service.

Note

To discover current settings, use the `networkadmin` command with the `-mib` function.

SYNTAX

```
networkadmin -mib [-enable] [-disable] [-community string] [-o]
```

OPTIONS

- `-enable`
Enables the SNMP services.
 - `-disable`
Disables the SNMP services.
 - `-community newstring`
Sets the SNMP community string to *newstring*.
-

Note

The `-community` switch is supported only with ManagementServer release 28 or later.

- `-o`
Executes the command without prompting for confirmation.

EXAMPLE

This command disables the SNMP services on SP A.

```
naviseccli -h ssl_spa networkadmin -mib -disable
```

```
Disable processing of SNMP read requests on SP A (y/n) [n]
```

OUTPUT

```
Storage Processor: SP A
SNMP MIB Status: Disabled
Community: public
```

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

networkadmin -route

Sets up the static routes.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli networkadmin` command with the `-route` function is used to set up the static routes to the specific hosts or networks.

SYNTAX

```
networkadmin -route [-sp a|b] [-portid portid -vportid vportid]
[-add -destination destination [-subnetmask subnetmask |
-prefixlength length] [-metric metric] [-persistent]]
[-delete -destination destination [-subnetmask subnetmask |
-prefixlength length]] [-list] [-o]
```

OPTIONS

`-sp a|b`

Specifies the destination SP. The default is the connected SP.

`-portid portid`

The `-portid` and the `-vportid` together specify the interface associated with the route. If you do not specify the interface, the system determines it.

`-vportid vportid`

The `-portid` and the `-vportid` together specify the interface associated with the route. If you do not specify the interface, the system determines it.

`-add`

Adds a network route.

`-destination destination`

Specifies the IPv4/IPv6 address or the hostname of the destination.

`-subnetmask netmask`

Specifies the IPv4 subnet mask value for the route entry. The default value is 255.255.255.255. It is valid only for an IPv4 destination address.

`-prefixlength length`

Specifies the prefix length and the part of the address to compare while determining the route. The default value is 128. It is valid only for an IPv6 destination address.

`-metric metric`

Determines the best route using the routing algorithm. The route with the lowest metric is preferred over another route.

`-persistent`

Specifies whether the route should persist across restarts.

`-delete`

Deletes an existing route.

`-list`

Displays the routes for the SP. By default, it displays the routes of the connected SP. It also displays the routes for a specific port/virtual port.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -h kirk-spa networkadmin -route -list
```

OUTPUT

```

SP A, Port 0, Virtual Port 0:
Protocol:                IPv4
Subnet Mask:             255.255.255.0
Gateway:                10.5.2.1
Originator:             OS
Metric:                 10
Destination Address:    10.5.2.0
Persistent:             YES
Static:                 NO

Protocol:                IPv4
Subnet Mask:             255.255.255.255
Gateway:                10.5.2.1
Originator:             OS
Metric:                 10
Destination Address:    10.255.255.255
Persistent:             YES
Static:                 NO

Protocol:                IPv4
Subnet Mask:             240.0.0.0
Gateway:                10.5.2.1
Originator:             OS
Metric:                 10
Destination Address:    224.0.0.0
Persistent:             YES
Static:                 NO

Protocol:                IPv4
Subnet Mask:             255.255.255.255
Gateway:                10.5.2.1
Originator:             OS
Metric:                 1
Destination Address:    255.255.255.255
Persistent:             YES
Static:                 NO

Protocol:                IPv6
Destination Prefix Length: 64
Originator:             OS
Metric:                 8
Destination Address:    3ffe:80c0:22c:47::
Persistent:             YES
Static:                 NO

```

networkadmin -set

Updates network information for an SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli networkadmin` command with the `-set` function changes one or more network settings for an SP. This command supports IPv6 network communication to the management ports of systems with VNX OE for Block version 04.28 or later. If you do not specify the IP address (either IPv4 or IPv6), then by default IPv4 information is displayed. The default IP address is IPv4. If both the `-ipv4` and `-ipv6` switches are specified, the information is displayed for both IP addresses, with the IPv4 information listed first.

Issue the command to the SP whose information you want to update.

If you omit the `-o` (override) option, the CLI displays a message in the following form to confirm the update.

Changing the name of SP <A|B> from oldname to newname (y/n) [n]?

Changing the IP address of SP <A|B> from old-address to new-address (y/n) [n]?

Changing the sub-net mask of SP <A|B> from old-mask to new-mask (y/n) [n]?

Changing the gateway address of SP <A|B> from old-gateway-address to new-gateway-address (y/n) [n]?

To make the change, enter **y**; otherwise, enter **n**.

Note

To discover current settings, use the `networkadmin` command with the `-get` function.

EMC service personnel set the network properties initially to work at your site. Do not change any value unless you are moving the SP to another LAN or subnet.

If you change any value, after you confirm, the SP restarts and uses the new value.

SYNTAX

```
networkadmin -set [-o] [-sp a|b] [-name name] [-portid portid
-vportid vportid] [-vlanid vlanid | -vlandisable]
[-ipv4 [-address address] [-subnetmask subnetmask] [-gateway
gateway-IP-address]]
[-ipv6 [-automatic | -disable | -manual
[[-globalprefix prefix | -address IPv6-address]
[-gateway gateway-IP-address]]]
[-speed [-duplex]]
```

OPTIONS

`-o`

Executes the command without prompting for confirmation.

`-sp [a|b]`

Modifies the properties of the specified SP. The default is the connected SP.

`-name name`

Changes the SP's network name to *name*. The maximum length of the name is 64 characters.

`-portid portid`

Specifies the management port ID. The default management port is 0.

`-vportid vportid`

Specifies the virtual port ID associated with the management port. The default virtual port is 0.

`-vlanid vlanid`

Specifies the VLAN ID to be used for the management port/virtual port. If you specify an invalid VLAN ID or if the ID is already in use, an error message appears displaying the valid range.

`-vlandisable`

Disables VLAN tagging on the virtual port.

`-ipv4`

Specifies the settings for IPv4.

- address *IP-address*
Changes the SP's IPv4 network address to *IP-address*.
- subnetmask *mask*
Changes the SP's IPv4 subnet mask to *mask*.
- gateway *gateway-IP-address*
Changes the SP's IPv4 or IPv6 gateway IP address to *gateway-IP-address*.
- ipv6
Specifies the settings for IPv6.
[-automatic | -disable | -manual]
Specifies whether you want to enable IPv6 on the SP automatically, disable IPv6 on the SP, or enable IPv6 on the SP manually. If you choose *manual*, also choose whether to set either a global prefix or a specific IPv6 address, and a gateway.
- globalprefix *prefix*
Changes the SP's IPv6 global prefix to *prefix*. If you are using this option, you cannot also use the -address *IPv6-address* option.
- address *IPv6-address*
Sets an IPv6 address on the SP, overriding any previous IPv6 configuration. If you are using this option, you cannot also use the -globalprefix *prefix* option.
- speed [-duplex] (CX4 series only)
Changes the speed on the management port for the target SP.

EXAMPLE #1

For SP A, this command changes the gateway address.

```
naviseccli -h ssl_spa networkadmin -set -gatewayaddress
```

```
Changing the gateway address of SP A from 123.456.789.011 to
123.456.789.011
```

```
This operation will cause a management server restart!
Do you wish to continue? (y/n) [n] y
```

EXAMPLE #2

This command enables IPv6 (if not enabled) and maintains the current settings.

```
naviseccli -address 10.5.2.12 -user a -password a -scope 0
networkadmin -set -ipv6 -manual
```

EXAMPLE #3

This command enables IPv6 (if not enabled) and sets the IPv6 global prefix.

```
naviseccli -address 10.5.2.12 -user a -password a -scope 0
networkadmin -set -ipv6 -manual -globalprefix 3ffe:80c0:22c:4c:
```

EXAMPLE #4

For SP A, this command changes the management port speed to 1000 Mbps/half duplex.

```
naviseccli -h ssl_spa networkadmin -set -speed 1000
-duplex half
```

```
Change the management port speed for SP A to 1000 Mbps/half duplex
DO YOU WISH TO CONTINUE? (y/n)?
```

EXAMPLE #5

For SP A, this command manually sets an IPv6 address.

```
naviseccli -h ss2_spa -user b -password b -scope 0
networkadmin -set -ipv6 -manual -address 3ffe:80c0:22c:4c:1319:8a2e:
370:7348
```

```
Setting the IPv6 Mode to Manual on SP A, Port 0, Virtual Port 0
This operation will cause a management server restart!
Do you wish to continue? (y/n)
```

OUTPUT

See the description section of this command. If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
NETADMIN_NOT_FOUND
NETADMIN_INVALID_NAME
Invalid command line parameters
Invalid speed request
```

port

Lists SP port information, removes an HBA entry from an SP's initiator list, and provides diagnostic capabilities.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli port` command lists SP port information, removes an HBA entry from an SP's permitted initiator list (initiator record), and provides diagnostic capabilities for SAN Copy initiator records and duplicate host initiator records. See also the `storagegroup -setpath` command in [storagegroup on page 168](#) and the `alpha` command in [alpha -set on page 27](#).

Note

To register an individual HBA with a system, you need to use `Unisphere`.

The optional `port -diagnose` switches are:

```
[-sancopy [-clean]] [-host]
```

SYNTAX

```
port -list [-hba|-arrayhba|-sp|-gname groupname |-uid UID]
[-all] [-arraycommpath] [-failovermode] [-unitserialnumber]
[-mac] [-initiatorcount] [-sfpstate] [-reads] [-writes]
[-bread] [-bwrite] [-qfull] [-physical] [-usage] [-prominfo]
-removeHBA [-o] [-hbauid HBAUID |-host host|-hostip ip|-all]
-diagnose [-o] -sancopy [-clean] |-host
```

OPTIONS

`-sancopy [-clean]`

Lists the SAN Copy initiator records where the initiator is the same as that of the targeted SP. If you include the optional `-clean` switch, the CLI prompts you to remove the initiator records that are displayed.

- host
Scans the host initiator records and displays duplicate hosts with the same host ID.
Without arguments, `-list` displays information about all ports in the system but does not list information given by `-all`, `-arraycommpath`, and `-failovermode`. This allows `-list` without an argument to produce output compatible with scripts written for old revisions of Secure CLI. Example 1 shows output without switches. Example 2 shows output with the `-all` switch. You can use one of the following optional switches with `-list`:
- all
Lists all port information, including any information that is new with this Unisphere release.
- arraycommpath
Displays the status of the `arraycommpath` setting: 0 if `arraycommpath` is disabled; 1 if `arraycommpath` is enabled.
- arrayhba
Displays information about all systems attached to this HBA, or about systems that were logged in at some point for which initiator records still exist. See Example 2 for a display.
- bread
Displays the number of blocks read for each port.
- failovermode
Displays the failover mode: 0, 1, 2, 3, or 4. The `storagegroup -failovermode` description provides more information.
- gname *groupname*
Specifies a storage group. *goupname* specifies the user-defined storage group name. If you omit the `-gname` and the `-uid` switches, the CLI describes all HBA ports connected to this storage group.
- hba
Returns information on HBA ports. The initiator and target `-hba` field returns different information depending on whether you are targeting a Fibre Channel or FCoE or an iSCSI HBA. When you target a Fibre Channel or FCoE HBA, the `-hba` returns the World Wide Name such as: 10:00:00:00:C9:20:DE:0C: 10:00:00:00:C9:20:DE:0C. When you target an iSCSI HBA, `-hba` returns the initiator name.
- initiatorcount
Displays the initiator information for each port on a given SP. Lists the number of initiators that are defined, logged in and not logged in.
- mac
Displays the MAC address for the front-end ports of an iSCSI system and also FCoE ports. Running the `port . . . -mac` command returns a *MAC Address* field for each port. When you run the `port . . . -mac` command, if any ports are Fibre Channel ports, the *MAC Address* field for those ports will display *Not Applicable*.
- physical (CX4 series only)
Displays the physical location of the port including I/O module and physical port information.
- reads
Displays the number of reads for each port.
- sfpstate
Displays the small form factor pluggable (SFP) state for each port. This lets you determine if an incompatible or faulty SFP is detected, in which the status is set to faulted.
- sp
Specifies SP ports only.

- uid *UID*
Specifies the storage group unique ID (also known as the World Wide Name, WWN). If you omit the `-gname` and the `-uid` switches, the CLI describes all HBA ports connected to this storage group.
- usage (CX4 series only)
Indicates how the port is used. For example, when either of the MirrorView enablers is installed, it displays MirrorView.
- prominfo
Displays the new EEPROM information for SFP+ SFPs inserted into the front-end ports.
- unitserialnumber
Specifies the `unitserialnumber` mode.
- writes
Displays the number of writes for each port.

The optional `port -removeHBA` switches are:

- hbaid *HBAUID*
Removes a specific initiator record specified with `-hbaid`. *HBAUID* specifies a unique host bus adapter identification number (must use uppercase).
- host *host*
Removes all initiator records that are logged out and registered, for the specified hostname.
- hostip *ip*
Removes all initiator records that are logged out and registered, for the specified host IP address. You can specify the IP address either as IPv4 or IPv6 address. If an incorrect IP address is specified, an error message appears as follows:

The IP address must be specified as an IPv4 address in dotted-decimal format or as an IPv6 address in colon-hexadecimal notation.

Note

If an IPv6 address is specified, EMC recommends that you use a global unicast address.

- all
Removes all initiator records that are logged out and registered, for the specified system.
- o
Executes the command without prompting for confirmation.

EXAMPLE #1

The following example shows sample output of the `port -list` command targeted to Fibre Channel HBAs and SPs:

Note

The **StorageGroup Name** field displays `None` if the HBA is not explicitly mapped to a user-defined, a non-user defined, a null, a management, or a physical storage group.

```
naviseccli -h ssl_spa port -list
```

```
Information about each HBA:
```

```
HBA UID:
10:10:10:10:10:10:10:10:10:10:10:10:10:10:10
```



```

Server Name:          matt
Server IP Address:    1.2.3.4
HBA Model Description: ModelABC
HBA Vendor Description: VendorABC
HBA Device Driver Name: N/A
Information about each port of this HBA:
SP Name:             SP A
SP Port ID:          0
HBA Devicename:     N/A
Trusted:             NO
Logged In:           NO
Defined:             YES
Initiator Type:      3
StorageGroup Name:  Private

Information about each HBA:

HBA UID:
20:00:00:00:C9:21:FE:D5:10:00:00:00:C9:21:FE:D5
Server Name:
20:00:00:00:C9:21:FE:D5:10:00:00:00:C9:21:FE:D5
Server IP Address:    UNKNOWN
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:
Information about each port of this HBA:
SP Name:             SP B
SP Port ID:          0
HBA Devicename:
Trusted:             NO
Logged In:           YES
Source ID:            7280384
Defined:             NO
Initiator Type:      128
StorageGroup Name:  None

Information about each HBA:

HBA UID:
AA:01:34:11:22:33:44:44:72:68:01:38:01:21:42:53
Server Name:          dochost.example.com
Server IP Address:    10.10.10.1
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:
Information about each port of this HBA:
SP Name:             SP A
SP Port ID:          0
HBA Devicename:
Trusted:             NO
Logged In:           NO
Defined:             YES
Initiator Type:      3
StorageGroup Name:  Documentation
SP Name:             SP A
SP Port ID:          1
HBA Devicename:
Trusted:             NO
Logged In:           NO
Defined:             YES
Initiator Type:      3
StorageGroup Name:  Documentation

Information about each HBA:

HBA UID:
20:00:00:00:C9:2B:53:D3:10:00:00:00:C9:2B:53:D3
Server Name:

```

```

20:00:00:00:C9:2B:53:D3:10:00:00:00:C9:2B:53:D3
Server IP Address: UNKNOWN
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:
Information about each port of this HBA:
SP Name:          SP A
SP Port ID:       0
HBA Devicename:
Trusted:         NO
Logged In:       YES
Source ID:       1315328
Defined:         NO
Initiator Type:  128

Information about each HBA:

HBA UID:
00:11:22:33:44:55:66:77:88:99:AA:BB:CC:DD:EE:FF
Server Name: abc.example.com
Server IP Address: 10.10.10.2
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:
Information about each port of this HBA:
SP Name:          SP A
SP Port ID:       0
HBA Devicename:
Trusted:         NO
Logged In:       NO
Defined:         YES
Initiator Type:  3
StorageGroup Name: PublicStorageGroup

Information about each SPPORT:

SP Name:          SP A
SP Port ID:       0
SP UID:          50:06:01:60:BC:E0:0C:12:50:06:01:63:3C:E0:0C:12
Link Status:     Up
Port Status:     Online
Switch Present:  YES
Switch UID:      10:00:00:05:1E:0E:E8:7B:20:1D:00:05:1E:0E:E8:7B
SP Source ID:    859648
I/O Module Slot: Onboard
Physical Port ID: 2

SP Name:          SP A
SP Port ID:       1
SP UID:          50:06:01:60:BC:E0:0C:12:50:06:01:63:3C:E0:0C:12
Link Status:     Up
Port Status:     Online
Switch Present:  YES
Switch UID:      10:00:00:05:1E:0E:E8:7B:20:1D:00:05:1E:0E:E8:7B
SP Source ID:    859648
I/O Module Slot: Onboard
Physical Port ID: 3

SP Name:          SP A
SP Port ID:       2
SP UID:          50:06:01:60:BC:E0:0C:12:50:06:01:63:3C:E0:0C:12
Link Status:     Up
Port Status:     Online
Switch Present:  YES
Switch UID:      10:00:00:05:1E:0E:E8:7B:20:1D:00:05:1E:0E:E8:7B
SP Source ID:    859648
I/O Module Slot: Onboard
Physical Port ID: 4

```

```

SP Name:          SP A
SP Port ID:       3
SP UID:           50:06:01:60:BC:E0:0C:12:50:06:01:68:3C:E0:0C:12
Link Status:      Up
Port Status:      Online
Switch Present:   YES
Switch UID:       10:00:00:05:1E:0E:E8:7B:20:1D:00:05:1E:0E:E8:7B
SP Source ID:     859648
I/O Module Slot:  Onboard
Physical Port ID: 5

```

EXAMPLE #2

The following example shows sample output of the `port -list -all` command targeted to Fibre Channel HBAs and SPs.

```
naviseccli -h ssl_spa port -list -all
```

```

Total number of initiators: 1

SP Name:          SP A
SP Port ID:       1
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP A
SP Port ID:       0
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP B
SP Port ID:       1
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP B
SP Port ID:       0
Registered Initiators: 1
Logged-In Initiators: 1
Not Logged-In Initiators: 0

SP Name:          SP A
SP Port ID:       2
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP A
SP Port ID:       3
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP B
SP Port ID:       2
Registered Initiators: 0
Logged-In Initiators: 0
Not Logged-In Initiators: 0

SP Name:          SP B
SP Port ID:       3
Registered Initiators: 0
Logged-In Initiators: 0

```

```

Not Logged-In Initiators: 0

Information about each HBA:
HBA UID:
20:00:00:00:C9:5E:B8:FF:10:00:00:00:C9:5E:B8:FF
Server Name:          nlpc46123.us.dg.com
Server IP Address:    10.14.46.123
HBA Model Description: Emulex LightPulse HBA - Storport
Miniport Driver
HBA Vendor Description: Emulex
HBA Device Driver Name: elxstor
Information about each port of this HBA:

SP Name:              SP B
SP Port ID:           0
HBA Devicename:       \\.\SCSI3:0:0:0
Trusted:              NO
Logged In:             YES
Source ID:             8327680
Defined:              YES
Initiator Type:       3
StorageGroup Name:    None
ArrayCommPath:        1
Failover mode:        1
Unit serial number:   Array

```

EXAMPLE #3

The following example shows sample output of the `port -list -hba` command targeted to an iSCSI HBA.

```
naviseccli -h IPAddress_SP port -list -hba
```

```

Information about each HBA:
HBA UID: iqn.1991-05.com.microsoft:nlpc20234
Server Name:          nlpc20234
Server IP Address:    25.24.23.235
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:
Information about each port of this HBA:

SP Name:              SP A
SP Port ID:           0
HBA Devicename:       \\.\SCSI2:0:0:1
Trusted:              NO
Logged In:             YES
Source ID:             4294967295
Defined:              YES
Initiator Type:       3
StorageGroup Name:    StorageGrpnlpc20234

Information about each SPPORT:

SP Name:              SP A
SP Port ID:           0
SP UID:               iqn.1992-04.com.emc:cx.apm00034901526.a0
Link Status:          Up
Port Status:          Online
Switch Present:       Not Applicable

SP Name:              SP A
SP Port ID:           1
SP UID:               iqn.1992-04.com.emc:cx.apm00034901526.a1
Link Status:          Down
Port Status:          Online
Switch Present:       Not Applicable

```

```

SP Name:          SP B
SP Port ID:       0
SP UID:           iqn.1992-04.com.emc:cx.apm00034901526.b0
Link Status:      Down
Port Status:      Online
Switch Present:   Not Applicable

SP Name:          SP B
SP Port ID:       1
SP UID:           iqn.1992-04.com.emc:cx.apm00034901526.b1
Link Status:      Down
Port Status:      Online
Switch Present:   Not Applicable

```

EXAMPLE #4

The following example shows sample output of the `port -list -mac` command targeted to an iSCSI system.

```
naviseccli -h 10.14.80.107 port -list -sp -mac
```

```
Information about each SPPORT:
```

```

SP Name:          SP B
SP Port ID:       0
SP UID:           iqn.1992-04.com.emc:cx.apm00034901792.b0
Link Status:      Down
Port Status:      Online
Switch Present:   Not Applicable
MAC Address:      00:60:16:01:83:AF

SP Name:          SP B
SP Port ID:       1
SP UID:           iqn.1992-04.com.emc:cx.apm00034901792.b1
Link Status:      Up
Port Status:      Online
Switch Present:   Not Applicable
MAC Address:      00:60:16:01:83:B1

SP Name:          SP A
SP Port ID:       0
SP UID:           iqn.1992-04.com.emc:cx.apm00034901792.a0
Link Status:      Down
Port Status:      Online
Switch Present:   Not Applicable
MAC Address:      00:60:16:01:82:89

SP Name:          SP A
SP Port ID:       1
SP UID:           iqn.1992-04.com.emc:cx.apm00034901792.a1
Link Status:      Up
Port Status:      Online
Switch Present:   Not Applicable
MAC Address:      00:60:16:01:82:8B

```

EXAMPLE #5

The following example shows sample output of the `port -list -mac` command targeted to a Fibre Channel system.

```
naviseccli -h 10.14.5.202 port -list -sp -mac
```

```
Information about each SPPORT:
```

```

SP Name:          SP A
SP Port ID:       1

```

```

SP UID: 50:06:01:60:90:21:26:3D:50:06:01:61:10:21:26:3D
Link Status: Up
Port Status: Online
Switch Present: NO
MAC Address: Not Applicable

SP Name: SP A
SP Port ID: 0
SP UID: 50:06:01:60:90:21:26:3D:50:06:01:60:10:21:26:3D
Link Status: Up
Port Status: Online
Switch Present: YES
Switch UID:
10:00:08:00:88:60:8E:77:20:0D:08:00:88:60:8E:77
SP Source ID: 7998739
MAC Address: Not Applicable

SP Name: SP B
SP Port ID: 1
SP UID: 50:06:01:60:90:21:26:3D:50:06:01:69:10:21:26:3D
Link Status: Up
Port Status: Online
Switch Present: NO
MAC Address: Not Applicable

SP Name: SP B
SP Port ID: 0
SP UID: 50:06:01:60:90:21:26:3D:50:06:01:68:10:21:26:3D
Link Status: Up
Port Status: Online
Switch Present: YES
Switch UID:
10:00:08:00:88:60:8E:77:20:0C:08:00:88:60:8E:77
SP Source ID: 7998483
MAC Address: Not Applicable

```

EXAMPLE #6

The following example shows sample output of the `port -list -mac` command targeted to a Fibre Channel over Ethernet (FCoE) system that has FC(0,1,2,3), 1G iSCSI(4,5), 10G iSCSI(6,7) and 10G FCoE(8,9) ports.

```
naviseccli -h peregrine-spa port -list -mac
```

```

Information about each SPPORT:
SP Name: SP B
SP Port ID: 3
MAC Address: Not Applicable

SP Name: SP B
SP Port ID: 2
MAC Address: Not Applicable

SP Name: SP B
SP Port ID: 0
MAC Address: Not Applicable

SP Name: SP B
SP Port ID: 1
MAC Address: Not Applicable

SP Name: SP A
SP Port ID: 3
MAC Address: Not Applicable

SP Name: SP A
SP Port ID: 2

```

```

MAC Address:          Not Applicable

SP Name:              SP A
SP Port ID:           0
MAC Address:          Not Applicable

SP Name:              SP A
SP Port ID:           1
MAC Address:          Not Applicable

SP Name:              SP B
SP Port ID:           4
MAC Address:          00:60:16:1A:3B:C2

SP Name:              SP B
SP Port ID:           5
MAC Address:          00:00:00:00:00:00

SP Name:              SP B
SP Port ID:           6
MAC Address:          00:60:16:32:12:1C

SP Name:              SP B
SP Port ID:           7
MAC Address:          00:00:00:00:00:00

SP Name:              SP A
SP Port ID:           4
MAC Address:          00:60:16:1A:36:52

SP Name:              SP A
SP Port ID:           5
MAC Address:          00:00:00:00:00:00

SP Name:              SP A
SP Port ID:           6
MAC Address:          00:60:16:32:0E:8E

SP Name:              SP A
SP Port ID:           7
MAC Address:          00:00:00:00:00:00

SP Name:              SP B
SP Port ID:           8
MAC Address:          00:60:16:3B:47:B2

SP Name:              SP B
SP Port ID:           9
MAC Address:          00:60:16:3B:47:B3

SP Name:              SP A
SP Port ID:           8
MAC Address:          00:60:16:3B:55:42

SP Name:              SP A
SP Port ID:           9
MAC Address:          00:60:16:3B:55:43

```

EXAMPLE #7

The following example shows sample output of the `port -list -sp -all` command.

```

SP Name:              SP B
SP Port ID:           0
SP UID:              50:06:01:60:B6:60:02:DC:50:06:01:68:36:60:02:DC

```

```

Link Status:                               Up
Port Status:                               Online
Switch Present:                            YES
Switch UID:
10:00:00:05:1E:E0:0B:E0:20:03:00:05:1E:E0:0B:E0
SP Source ID:                              932096
ALPA Value:                                0
Speed Value :                              8Gbps
Auto Negotiable :                          NO
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value:                           Auto
MAC Address:                               Not Applicable
SFP State:                                 Online
Reads:                                     0
Writes:                                    0
Blocks Read:                               0
Blocks Written:                            0
Queue Full/Busy:                           0
I/O Module Slot:                           0
Physical Port ID:                           0
Usage:                                     General
SFP/Connector EMC Part Number: 019-078-042
SFP/Connector EMC Serial Number: 000000000000000
SFP/Connector Vendor Part Number: N/A
SFP/Connector Vendor Serial Number: AGL1211D3004397
SFP/Connector Supported Speeds:
2Gbps
4Gbps
8Gbps

```

EXAMPLE #8

The following example shows sample output of the port -list -sfpsstate command targeted to FCoE ports, 10G iSCSI, and FC ports.

```

SPF Normal:

SP Name:          SP A
SP Port ID:      11
SP UID:          50:06:01:60:16:21:02:B6:50:06:01:60:16:25:02:B6
Link Status:     Down
Port Status:     DISABLED
Switch Present:  NO
SFP State:       Online

SPF Faulted:

SP Name:          SP A
SP Port ID:      8
SP UID:          50:06:01:60:BB:20:13:0D:50:06:01:68:3B:24:13:0D
Link Status:     Down
Port Status:     DISABLED
Switch Present:  NO
SFP State:       Faulted

```

EXAMPLE #9

The following example shows the output for the port -list -prominfo command.

```

Information about each SPPORT:

SP Name:          SP A
SP Port ID:      2

```



```

SFP/Connector EMC Part Number:      TBD
SFP/Connector EMC Serial Number:    TBD
SFP/Connector Vendor Part Number:   TBD
SFP/Connector Vendor Serial Number: TBD
SFP/Connector Supported Speeds:
2Gbps
4Gbps
8Gbps

SP Name: SP A
SP Port ID: 3
SFP/Connector EMC Part Number:      TBD
SFP/Connector EMC Serial Number:    TBD
SFP/Connector Vendor Part Number:   TBD
SFP/Connector Vendor Serial Number: TBD
SFP/Connector Supported Speeds:
2Gbps
4Gbps
8Gbps
...

```

powersaving

Enables or disables the power saving settings.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli powersaving` command allows you to enable or disable power saving settings on a system. It displays the system's idle time and the status of the power saving settings.

SYNTAX

```

powersaving -globalsettings [-o] [on|off]
-info [-systemidletime] [-globalpowersavingsettings]

```

OPTIONS

`-globalsettings`

Enables or disables power saving settings on a system. The system sends a confirmation message only when you specify the `-o` switch.

The following message is displayed, if the settings are turned on:

```

Turning global power saving settings on will spin down all
eligible unused disks and all power saving settings on
Storage Pools will take effect. Do you want to proceed?

```

The following message is displayed, if the settings are turned off:

```

Turning global power saving settings off will allow all
drives spinning up to normal. Do you want to proceed?

```

`-info [-systemidletime] [-globalpowersavingsettings]`

Displays the system idle time and the status of the global power saving settings.

EXAMPLE

```

naviseccli -h IP address powersaving -info

```

```

System Idle Time: 15 min
Global Power Saving Settings: On

```

rebootSP

Reboots an SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command reboots one or both SPs. If you omit switches, the current SP is rebooted.

Note

If `rebootSP` is executed while I/O processes are active, I/O errors may occur and/or the Unisphere agent process may hang. If the agent process hangs, you must reboot the host on which the Unisphere agent is running. If you are binding LUNs assigned to one SP, you should not reboot the peer SP until the binding process has completed. Otherwise, until the binding process is completed, you will see error messages each time the SP is polled.

Note

See the `rebootpeerSP` command, to reboot the peer SP of the target SP.

SYNTAX

```
rebootSP [-both] [-o]
```

OPTIONS

`-both`
Reboot both SPs.

Note

If write cache is enabled, it is not allowed to reboot both SPs unless you disable write cache first or use `-o` option to override the warning and reboot directly.

`-o`
Executes the command without prompting for confirmation.

EXAMPLE #1

```
naviseccli -h ss1_spa rebootSP
```

```
WARNING: This operation will reboot the target SP.
Verify that the SPS on the surviving SP is operating correctly.
Otherwise write cache must be disabled:
naviseccli -h IPAddress cache -sp -modify -wc off.
Failure to disable write cache may result in the loss of data.
DO YOU WISH TO CONTINUE? (y/n)
```

EXAMPLE #2

```
naviseccli -h ss1_spa rebootSP -both
```

```
Error: rebootSP command failed
Before you can execute the rebootSP command, you need to either
disable write cache or override the warning.
To disable write cache, run the "naviseccli -h IPAddress cache -sp
-modify -wc off" command.
To override the warning and reboot both SPs directly, run the
"rebootSP -both -o" command.
```

OUTPUT

There is no output. Errors are printed to `stderr`.

rebootpeerSP

Reboots a peer SP.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Lets you reboot the peer SP of the target SP in a dual-SP system. For the `rebootpeerSP` command, the peer SP is the SP on the system to which the command is not targeted.

Note

See the `rebootSP` command, to reboot the target SP or both SPs on the system.

SYNTAX

```
rebootpeerSP [-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation.

EXAMPLE

The following example reboots the peer SP of the target SP. The target SP is `ss1_spa`; the peer SP is `ss1_spb`.

```
naviseccli -h ss1_spa rebootpeerSP
```

```
WARNING: This operation will reboot the target SP.
Verify that the SPS on the surviving SP is operating correctly.
Otherwise write cache must be disabled:
naviseccli -h IPAddress cache -sp -modify -wc off.
Failure to disable write cache may result in the loss of data.
DO YOU WISH TO CONTINUE? (y/n)
```

remoteconfig -getconfig

Displays information about a host or SP agent.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `remoteconfig` command displays and sets agent configuration information. An agent can be any host agent or SP agent. These commands let you get and set the configuration of the agent running on a remote host, scan all the devices in the system on a remote host, and stop the agent running on a host.

Note

If you are running VMware ESX Server and issuing this command to a virtual machine (VM), it cannot target the host agent because the Unisphere Host Agent is supported only on ESX Server.

The `naviseccli remoteconfig` command with `-getconfig` displays information on the agent running on an SP. To display information on the agent for the specified host, see [server -remoteconfig on page 289](#). If you omit switches, the command displays all agent information. You can request specific information with switches.

To write the agent information to a file, use the `-write filename` switch. If you use `-write`, the CLI writes all information to the file; you cannot select specific information with switches. If a file with the same name exists, the CLI will verify to overwrite.

Note

After listing agent information, you can change the agent settings with the `remoteconfig -setconfig` command.

SYNTAX

```
remoteconfig -getconfig [-description] [-contact] [-users]
[-interval] [-write filename]
```

OPTIONS

`-description`

Displays contents of the description field.

`-contact`

Displays contents of the contact field.

`-users`

Displays contents of the users field.

`-interval`

Displays contents of the polling interval field.

`-write filename`

Writes an image of the configuration file as it exists on the host or SP to file *filename*. The file is written in the same format as the `agent.config` file, which differs from the output of the `getconfig` function without the `-write` switch. You can use this file to set the configuration for other agents using the `remoteconfig setconfig -f filename` function. No other switch is allowed with `-write`.

EXAMPLE

This command lists all agent information.

```
naviseccli -h ssl_spa remoteconfig -getconfig
```

OUTPUT

See [remoteconfig -getconfig on page 147](#). If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
RAC_AGENT_NOT_RUNNING
```

remoteconfig -setconfig

Changes configuration information for a host or SP agent.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli remoteconfig` command with `-setconfig` changes the configuration information for the agent running on an SP. To change configuration

information for the agent running on the specified host, see [server -remoteconfig on page 289](#).

Note

The comments in the host agent configuration file are removed. When your changes to the configuration file take effect, comments that previously resided in the file are permanently lost.

The CLI prompts for confirmation of every change you specify unless you use the `-o` (override) or `-f` switch. To change the setting for the agent, enter `y`. The confirmation queries for all settings are as follows.

```
Change host description from old to new <y/n> [y]?
Change contact information from old to new <y/n> [y]?
Add user in the authorized user list <y/n>[y] ?
Remove user from the authorized user list <y/n>[y] ?
Change polling interval to new-interval <y/n>[y] ?
```

To have the CLI obtain the agent information from a file (perhaps a file written with the `getconfig -write` function), use the `-f filename` switch. Only the `-o` (override) switch is allowed with the `-f` switch.

You can use the `-getconfig` function to see how settings are formatted. Make sure that your entries are valid for the agent you are configuring. A wrong setting can make an agent inaccessible.

SYNTAX

```
remoteconfig -setconfig [-o] [-adduser userlist] [-contact
contactinfo] [-description description] [-f filename] [-interval
interval] [-rmuser userlist]
```

OPTIONS

- `-o`
Executes the command without prompting for confirmation.
- `-adduser userlist`
Adds the users specified in *userlist* to the authorized user list. The user name form is `username@hostname`. Separate the user entries with a space.
- `-contact contactinfo`
Specifies new contact information for the agent description field.
- `-description description`
Specifies new contents for the agent description field.
- `-f filename`
Reads the configuration information from file *filename*. This file must be in the correct format (same as the agent configuration file used in earlier Unisphere revisions). The `-getconfig -write` command automatically creates files in the correct format. When you use `-f`, the CLI does not prompt for confirmation; the `-o` switch is not needed.
- `-interval interval`
Changes the polling interval in seconds to *interval*. The valid range is 1 to 120. This switch may not apply to all system types.
- `-rmuser userlist`
Removes the specified users from the authorized user list.

EXAMPLE

This command adds the user `tsjones` on host `viking` to the agent privileged users list.

```
naviseccli -h ssl_spa remoteconfig -setconfig -adduser tsjones@viking
```

```
Add tsjones@viking in the authorized user list <y/n>[y]?
```

OUTPUT

See [remoteconfig -getconfig on page 147](#). If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
RAC_AGENT_NOT_RUNNING
RAC_REMOTECONFIG_NOT_SUPPORTED
INVALID_MEGAPOLL
INVALID_USEEXPLICITDEVNAME
RAC_FILE_NOT_FOUND
```

removerg

Removes a RAID group.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli removerg` command removes a specified RAID group.

Note

This command cannot remove thin pools.

SYNTAX

```
removerg rgID
```

OPTIONS

rgID

RAID group identification number. *rgID* specifies a number between 0 and the maximum number of RAID groups supported by the RAID group system.

EXAMPLE

This command removes RAID group 1.

```
naviseccli -h ssl_spa removerg 1
```

resetandhold

Resets an SP and holds the reboot following the power on self tests (POST).

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli resetandhold` command, resets an SP and holds the reboot following the power on self tests (POST), to allow an SP to be reset in a safe manner for systems where pulling and reseating the SP is not permitted during operation. Because

the boot process is held, a subsequent `rebootpeerSP` command must be issued to clear the hold and reboot the SP, to return the SP to service.

Note

Issue `rebootpeerSP` after you issue `resetandhold`, when you are ready to return the SP to service.

SYNTAX

```
resetandhold
```

EXAMPLE

The following example resets and holds the reboot for the targeted SP on the system.

```
naviseccli -h ssl_spa resetandhold
```

```
This operation will cause the system to reset and
then hold the reboot process following the Power On Self
Tests (POST). To clear the hold state and complete the SP reboot,
a rebootpeersp command must be issued to the peer SP.
```

```
DO YOU WISH TO CONTINUE? (y/n)?
```

responsetest

Tests Event Monitor response.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli responstest` command creates an artificial event to test event monitor response.

The `responsetest` command is designed for use without the event monitor User Interface (UI). That is, you can use `responsetest` in conjunction with the event monitor template file, `navimon.cfg`, which you can edit to specify the event monitor responses you want. If you have used the event monitor UI to set up an event monitor template, do not use this command.

Event monitor functions accessed via `responsetest` are decentralized, providing email and pager support from local systems only. The EMC phone home feature is not available via this command.

SYNTAX

```
responsetest -email response-parameters | -emailpage response-
parameters | -modempage response-parameters |
-snmppage response-parameters
```

OPTIONS

response-parameters

Allows different parameters for each type of response you specify, as follows:

-email

-smtpmailserver *smtpmailserver_hostname*

Specifies the SMTP mail server; required.

-destaddress *destination_email_address*

Specifies the destination email address; required.

[-subject "*text*"]

The *text*, within quotation marks, specifies the subject of the email; optional.

`[-cc destination_email_address]`
 The email address to send a cc (copy) to; optional.

`[-message "text" -file filename]`
 The *text*, in quotation marks, describes the event; optional. Or you specify a file with descriptive text using `-file`.

`[-sender sender_email_address]`
 Specifies the sender address; optional. If omitted, CLI inserts the agent hostname as the *sender_email_address*.

`-emailpage`

`-smtpmailserver smtpmailserver_hostname`
 Specifies the SMTP mail server; required.

`-destaddress destination_email_address`
 Specifies the destination email address; required.

`[-subject "text"]`
 The *text*, within quotation marks, specifies the subject of the email; optional.

`[-cc destination_email_address]`
 The email address to send a cc (copy) to;

`[-message "text" -file filename]`
 The *text*, in quotation marks, describes the event; optional. Or you specify a file with descriptive text using `-file`.

`[-sender sender_email_address]`
 Specifies the sender address; optional. If omitted, CLI inserts the agent hostname as the *sender_email_address*.

`-modempage`
 (Not supported on AX4-5 series and AX series)

`-destnumber phone-number`
 Specifies the pager phone number (with area code) and is required.

`-msgnumber phone-number`
 Specifies the number that will appear on the pager display (with area code).

`[-comport number]`
 Specifies the *number* of the PC COM. If omitted, the CLI assumes COM 1; optional.

`[-messagedelay seconds]`
 Specifies the delay between sending the destination phone number to the modem and sending the message to the modem. This allows enough time for the receiving modem to answer the call. Optional.

`[-dialcommand command]`
 Specifies the modem dial command to be used by the agent. Optional.

`[-initcommand command-to-modem]`
 Specifies the modem dial command to be used by the agent. The `initcommand` allows for some custom modem configuration before the data is sent to it. Optional.

`-community community`
 Specifies a community on the snmp management host; optional.

`-desthost destination-hostname`
 Specifies an snmp management hostname; required.

`-device device`
 Specifies the device represented by the snmp management host. Required.

EXAMPLE

This command sends a modem page test event.

```
naviseccli -h ssl_spa responsetest modempage -destnumber -msgnumber
```

OUTPUT

Result of the test.

sc_off

Disables data access control.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli sc_off` command disables data access control and turns off shared storage features. You can use this command to reset the system to its factory unshared storage settings.

CAUTION

Be aware that executing `sc_off` resets the system to its factory unshared storage settings, permanently destroying all previously implemented shared storage features.

SYNTAX

```
sc_off [-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -h ssl_spa sc_off
```

OUTPUT

The first time you run the `sc_off` command, you will see an output similar to the following:

```
You have asked to remove all shared storage components
from the indicated system.

The following functions will be performed, if necessary:
Clear the system's password
Turn fairness off
Set the default storage group to ~physical
Remove all initiator records
Remove all user defined storage groups
Do you wish to proceed? (y/n)? y

Clearing the system's password
Old Password: ***
Turning fairness off
Setting the default storage group to ~physical
Removing all initiator records
Removing all user defined storage groups
```

If you run the command a second time, and the first time through completed successfully, the output will be as follows:

```
You have asked to remove all shared storage components
from the indicated system.

The following functions will be performed, if necessary:
Clear the system's password
Turn fairness off
Set the default storage group to ~physical
```

```
Remove all initiator records
Remove all user defined storage groups
Do you wish to proceed? (y/n)? y
```

setfeature

Adds a software feature to a LUN controlled by an SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setfeature` command adds and removes a software driver feature to a LUN owned by an SP.

SYNTAX

```
setfeature -on|-off -feature feature -lun lun| -lunuid uid
```

OPTIONS

`-on` or `-off`

Adds or removes the feature.

`-feature feature`

Identifies the feature; for example, RM for remote mirroring and sancopy for SAN Copy.

`-lun lun` or `-lunuid uid`

Specifies the LUN. You can use the LUN number or the LUN unique ID, UID (which is the World Wide Name, WWN).

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`. Other errors:

```
COMMAND_NOT_SUPPORTED
FEATURE_REQUIRED
LUN_AND_LUNUID_SPECIFIED
```

setsniffer

Starts a background verify process or changes verify process parameters.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setsniffer` command starts a background verify process or changes the parameters for the verify process. You can indicate a specific LUN, all LUNs in a RAID group, or all LUNs in a system.

Note

The `setsniffer` command does not support thin LUNs.

SYNTAX

```
setsniffer lun-number| [-rg rg-number] [-all] [-bv] [-cr]
```

OPTIONS**-all**

Applies sniffer parameters to all LUNs in the system. The target SP must own one LUN at minimum.

lun-number

Applies sniffer parameters to the specified LUN. *lun* specifies the logical unit number. You must target the SP that owns the specified LUN.

-rg *rg-number*

Applies sniffer parameters to all LUNs in the specified RAID group.

-bv

Starts a full unit background verify.

-cr

Clears background verify reports.

EXAMPLE

Starts a background verify on LUN 0.

```
naviseccli setsniffer 0 -bv
```

OUTPUT

There is no output. Errors are printed to `stderr`.

setsniffverify

Enables or disables sniff verify on all the disks on the array.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setsniffverify` command can enable or disable sniff verify on all the disks on the array, including bound and unbound region. This command with no option displays the current sniff verify state.

SYNTAX

```
setsniffverify [-on | -off]
```

OPTIONS**-on | off**

Enables or disables sniff verify on all disks.

-on

Enables sniff verify on all disks.

-off

Disables sniff verify on all disks.

EXAMPLE #1

This example enables the disk based sniff verify. There is no output. Errors are printed to `stderr`.

```
naviseccli setsniffverify -on
```

EXAMPLE #2

If no option is provided, the command will display the current sniff verify state.

```
naviseccli setsniffverify
sniffverify is ENABLED
```

setspstime

Sets SPS or BBU battery test time.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setspstime` command sets the SPS battery test time. If no switches are specified, the command returns the current SPS battery test time. The SPS test time is displayed for the client's time zone.

SYNTAX

```
setspstime [-d day-num] [-h hour-num] [-m minute-num]
[-nolocal]
```

OPTIONS

`-d day-num`

Sets the day of the week (0-6, 0=Sunday).

`-h hour-num`

Sets the hour of the day (0-23, 0=12:00 AM).

`-m minute-num`

Sets the minute of the hour (0-59).

Note

To set the SPS or BBU test time, all switches must be specified.

`-nolocal`

The SPS test time is displayed for the client's time zone.

Note

`-nolocal` is required if your system is communicating with a host agent version less than 6.5. You cannot use this switch with version 6.5 or higher.

EXAMPLE

This example sets the SPS battery test time to Sunday at 1:00 AM.

```
naviseccli -h ss1_spa setspstime -d 0 -h 1 -m 0
```

OUTPUT

There is no output. Errors are printed to `stderr`.

setsptime

Sets the SP system clocks.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setsptime` command sets the system time for both SPs.

SYNTAX

```
setsptime [-M month-num] [-D day-num] [-Y year-num]
[-h hour-num] [-m minute-num]
```

OPTIONS

- M *month-num*
Specifies the month (1-12).
- D *day-num*
Specifies the day (1-31).
- Y *year-num*
Specifies the year; use four digits.
- h *hour-num*
Specifies the hour (0-23).
- m *minute-num*
Specifies the minute (0-59).

EXAMPLE

This example sets the SP system clock to March 13, 2002, 9:00 am.

```
naviseccli -h data1_spa setsptime -M 3 -D 13 -Y 2002 -h 9
```

OUTPUT

There is no output. Errors are printed to `stderr`.

setstats

Sets statistics logging.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli setstats` command sets statistics logging. If no optional switches are present, the command returns the current state of statistics logging. The SP maintains a log of statistics for the LUNs, disk modules, and system caching that you can turn on and off. When enabled, logging affects system performance, so you may want to leave it disabled unless you have a reason to monitor performance.

The log uses a 32-bit counter to maintain the statistics numbers. When the counter is full, the statistics numbers restart at zero. As a result, you will see a sudden decrease in a statistics number if you view it shortly before the counter is full and shortly after the counter restarts at zero. If you want to keep the log turned on for more than two weeks, we recommend that you reset the log about every two weeks, so you know when the numbers start at zero.

The command affects only the SP you specify.

SYNTAX

```
setstats -on|-off [-powersavings] [-on|-off]
```

OPTIONS

- `setstats -on|-off`
Enables/disables the performance statistics logging.

Note

If statistics logging is off, certain options in other commands are disabled. The options for the other commands specify whether statistics logging must be on for their options to be valid.

-powersavings

Enables/disables the disk spin-down statistics logging. It affects only the disk spin-down statistics.

`-on|-off`

Returns the current state of power saving statistics logging.

EXAMPLE

The sample output is as follows:

```
naviseccli -h ss1_spa setstats -on
```

```
naviseccli -h ss1_spa setstats -powersavings
```

```
Disk Spin Down statistics logging is DISABLED
```

```
OR
```

```
Disk Spin Down statistics logging is ENABLED
```

```
Start Power Savings Log Timestamp: xx/xx/xxxx xx:xx:xx
```

OUTPUT

There is no output. Errors are printed to `stderr`.

shutdown

Shuts down and powers off the system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli shutdown` command shuts down and powers off the system you specify.

SYNTAX

```
shutdown [-o]
```

OPTIONS

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command shuts down and powers off the specified system.

```
naviseccli -h ss1_spa shutdown
```

```
Before you shutdown the system, stop all I/O from any servers connected to the system and save all unsaved data to disk. During the shutdown process you will lose access to all data on the LUNs but you will not lose any data. When you power up the system you will regain access to the LUNs.
```

```
Do you want to shutdown and power off the subsystem now
(y/n) ?
```

OUTPUT

None if the command succeeds; status or error information if it fails.

shutdownsp

Shuts down and powers off the current SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli shutdownsp` command shuts down the current storage processor (SP) of the system.

SYNTAX

```
shutdownsp [-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation. The confirmation message varies depending on the system.

EXAMPLE

The confirmation message for a VNX system is as follows:

```
WARNING: This operation will shutdown the SP and hold it in reset.
Verify that SPS on the surviving SP is operating correctly.
Otherwise write cache must be disabled: naviseccli -h IPAddress
cache -sp -modify -wc off.
To reconnect please reboot the SP - this can be accomplished by
using the
rebootpeerSP command.
Do you want to shutdown and hold the SP now? (y/n) ?
```

OUTPUT

None if the command succeeds; error information if used with any other system.

shutdownpeersp

Shuts down and powers off the peer SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli shutdownpeersp` command shuts down the peer storage processor (SP).

SYNTAX

```
shutdownpeersp [-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation.

The confirmation message varies depending on the system.

EXAMPLE

The confirmation message for a VNX system is as follows:

```
WARNING: This operation will shutdown the SP and hold it in reset.
Verify that SPS on the surviving SP is operating correctly.
Otherwise write cache must be disabled: naviseccli -h IPAddress
cache -sp -modify -wc off.
To reconnect please reboot the SP - this can be accomplished by
using the
rebootpeerSP command.
Do you want to shutdown and hold the SP now? (y/n) ?
```

OUTPUT

None if the command succeeds; error information if used with any other system.

spcollect -info

Displays the setting of the automated SPCollect service.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli spcollect` command with the `-info` function lets you display the setting of the automated SPCollect functionality. The setting appears as Enabled when the SPCollect automated functionality is turned on; Disabled when it is turned off.

The automated SPCollect service is part of the `Peerwatch` service included with VNX OE for Block version 02.19.xxx.5.yyy or higher.

SYNTAX

```
spcollect -info
```

EXAMPLE

This command shows you the setting of the automated SPCollect functionality for the specified system. The Enabled setting, indicating that the SPCollect process is turned on, appears in the output above.

```
naviseccli -address ss\1_spa spcollect -info
```

OUTPUT

```
AutoExecution: Enabled
```

spcollect -set

Sets the automated SPCollect service on or off.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli spcollect` command with the `-set` function lets you set the automated SPCollect process on or off.

The automated SPCollect service is part of the `Peerwatch` service included with VNX OE for Block version 02.19.xxx.5.yyy or higher.

Note

Use `sppcollect -info` to verify whether the service is enabled or disabled.

SYNTAX

```
sppcollect -set -auto on|off [-o]
```

OPTIONS

`-auto on|off`

Sets the automated SPCollect process on or off. If you indicate *on*, it turns the automation service on; if you indicate *off*, it turns the automation service off.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command turns the automated SPCollect process on.

```
naviseccli -address ssl_spa sppcollect -set -auto on
```

OUTPUT

None if the command succeeds; status or error information if it fails.

sppcollect

Selects and places a collection of system log files into a single .zip file.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli sppcollect` command selects a collection of system log files and places them in a single .zip file on the system. You can retrieve the file from the system using the `managefiles` command.

Note

This command is not related to the automated SPCollect service that is part of the Peerwatch service included with VNX OE for Block version 02.19.xxx.5.yyy or higher. This command lets you run the SPCollect functionality separate from the scheduled service. See the `sppcollect -set` command to enable or disable the automated SPCollect service.

NOTICE

The SPCollect functionality can affect system performance (may degrade system performance).

SYNTAX

```
sppcollect
```

EXAMPLE

This example creates a .zip file on the system, which contains system log files.

The .zip file has the following format:

```
arrayserialnumber_SPx_date_time_spsignature_data.zip
```

```
naviseccli -h ssl_spa sppcollect
```

arrayserialnumber is the system serial number.

x Is SP A or B.

date Is the date of collection.

time Is the time of collection.

spsignature Is the SP signature

OUTPUT

There is no output. Errors are printed to `stderr`.

sportspeed

Set or get the SP port speed.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli sportspeed` command gets or sets the SP port speed. You can set the port speed on either SP directly or through the other SP. The port speed you can set and the options available depend upon your system type.

For a port speed to work, the HBA, switch, and SP must all support it. You can set the ports of an SP to different speeds if the ports connect to switches or HBAs with different speeds.

CAUTION

Incompatible port speeds will prevent the two devices from communicating. The CLI prompts for confirmation as follows before executing the `-set` command.

```
Changing the port speed is a disruptive function. Verify that
the other device connected to the system, supports the new
speed. If the speed settings are incompatible, the system will
be unable to communicate with the other connected device. Are
you sure you want to change speed of SP x Port n from n Gigabit
to n Gigabit <y/n>?
```

SYNTAX

```
sportspeed -get [-isauto] [-avail] [-requested] [-type]
[-physical] [-sfp] [-all] -set -sp a|b
-portid portid speed [-o]
```

OPTIONS

`-isauto`

Indicates whether the Auto-Negotiate feature is turned on or off. The CLI returns a yes value if the feature is turned on, a no value if the feature is turned off, and a not supported value if the feature is not supported.

Note

If you have the Auto-Negotiate feature turned on, the Speed Value that is generated in the output does not have a value of auto. The value reflects the speed at which the port is currently operating. If you have the Auto-Negotiate feature turned on and you specify the `-requested` switch, the Requested Value displays a value of auto.

`-avail`

- Indicates the available speeds for a specified port.
- requested
 - Displays the speed requested through `spportspeed -set`. The speed that is requested may not equal the actual speed of the port. For example when the Auto-Negotiate feature is turned on, the Requested Value displays auto and the Speed Value displays the speed at which the port is currently operating.
- type
 - Displays the connection type (Fibre Channel or iSCSI) for each port along with the speeds.
- physical (CX series only)
 - Displays the physical location of the port.
- or
- sp *a/b*
 - Specifies the SP on which to set the port speed (*a/b*).
- portid *portid speed*
 - portid* specifies the port on the SP (0, 1, 2, or 3).
 - speed* specifies the port speed: See the note above. For Fibre Channel systems, port speed is 1 for 1 Gb mode, 2 for 2 Gb mode, 4 for 4 Gb mode, auto to initiate the Auto-Negotiate feature. For iSCSI systems, port speed is 10 for 10 Mb mode, 100 for 100 Mb mode, 1000 for 1000 Mb mode, auto to initiate the Auto-Negotiate feature.
- o
 - Executes the command without prompting for confirmation.

EXAMPLE #1

```
naviseccli -h ssl__spa spportspeed -get -physical
```

```
Storage Processor :      SP A
Port ID :             0
Speed Value :         8Gbps
Auto Negotiable :     Yes
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value:      Auto
Connection Type:      Fibre Channel
I/O Module Slot:      Onboard
Physical Port ID:     2
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor :      SP A
Port ID :             1
Speed Value :         8Gbps
Auto Negotiable :     Yes
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value:      Auto
Connection Type:      Fibre Channel
I/O Module Slot:      Onboard
Physical Port ID:     3
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps
```

```

Storage Processor :      SP A
Port ID :              2
Speed Value :          8Gbps
Auto Negotiable :      Yes
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value:       Auto
Connection Type:       Fibre Channel
I/O Module Slot:       Onboard
Physical Port ID:      4
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor :      SP A
Port ID :              3
Speed Value :          8Gbps
Auto Negotiable :      Yes
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value:       Auto
Connection Type:       Fibre Channel
I/O Module Slot:       Onboard
Physical Port ID:      5
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

```

EXAMPLE #2

naviseccli ssportspeed -get -sfp

```

Storage Processor : SP A
Port ID : 0
Speed Value : 8Gbps
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 1
Speed Value : 8Gbps
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 2
Speed Value : N/A
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 3

```

```

Speed Value : N/A
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 0
Speed Value : N/A
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 1
Speed Value : 8Gbps
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 2
Speed Value : N/A
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 3
Speed Value : N/A
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

```

EXAMPLE #3

```
naviseccli ssportspeed -get -all
```

```

Storage Processor : SP A
Port ID : 0
Speed Value : 8Gbps
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: Auto
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 0
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 1
Speed Value : 8Gbps
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps

```

```

8Gbps
Auto
Requested Value: Auto
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 1
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 2
Speed Value : N/A
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: 4Gbps
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 2
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP A
Port ID : 3
Speed Value : N/A
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: 4Gbps
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 3
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 0
Speed Value : N/A
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: 4Gbps
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 0
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 1
Speed Value : 8Gbps

```

```

Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: Auto
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 1
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 2
Speed Value : N/A
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: 4Gbps
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 2
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

Storage Processor : SP B
Port ID : 3
Speed Value : N/A
Auto Negotiable : No
Available Speeds:
2Gbps
4Gbps
8Gbps
Auto
Requested Value: 4Gbps
Connection Type: Fibre Channel
I/O Module Slot: 0
Physical Port ID: 3
SFP Supported Speeds:
2Gbps
4Gbps
8Gbps

```

EXAMPLE #4

```
navisecli -h ssl_spa ssportspeed -set -sp a -portid 0 2
```

Changing the port speed is a disruptive function. Verify that the other device connected to the system, supports the new speed. If the speed settings are incompatible, the system will be unable to communicate with the other connected device.

```
Are you sure you want to change speed of SP A Port n from
1 Gigabit to 2 Gigabit <y/n>? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Not supported error message is printed to stderr.

storagegroup

Creates and configures storage groups.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagegroup` command lets you create and manage shared systems, that is, systems that use storage groups and have the storage group option. For most command operations, the CLI prompts for confirmation unless you use the `-o` switch.

You can use the CLI command `storagegroup -connecthost` to assign a storage group to all HBAs in a host, or the command `storagegroup -setpath` to assign a storage group to a specific HBA in a host. To change the settings of a storage group that is already assigned, use the command `storagegroup -sethost`. To disconnect a host from all its storage groups, use `storagegroup -disconnecthost`.

The `storagegroup` command has many switches, some of which have functions that you also indicate with switches.

SYNTAX

```
storagegroup [-addhlu switches] |
[-addsnapshot switches] |
[-chgname switches] |
[-connecthost switches] |
[-create switch] |
[-destroy switch] |
[-disconnecthost switch] |
[-enable] |
[-list switches] |
[-removehlu switches] |
[-removesnapshot switches] |
[-sethost switches] |
[-setpath switches] |
[-shareable switch] |
```

OPTIONS

`storagegroup` command, `-addhlu` *switch* `-addhlu -gname`

`StorageGroupName` | `-uid` *StorageGroupUID* `-hlu` *hlunumber* `-alu` *alunumber*

Adds a host LUN to a storage group. This command requires a bound LUN. If required, use the `bind` command to create a LUN. You must use `-gname` or `-uid` and two other switches with `-addhlu` as follows. Also, see the example in [unbind on page 181](#).

`-gname` *StorageGroupName*

Specifies the name of the storage group.

`-uid` *StorageGroupUID*

Specifies the storage group unique identification number.

- For Fibre Channel systems, the UID uses the World Wide Name (WWN) format such as: 10:00:00:00:C9:20:DE:0C:10:00:00:00:C9:20:DE:0C.
- For iSCSI systems, the UID uses the initiator name which uses the IQN format such as: iqn.1992-04.com.emc:cx.apm00034901792.b0.

-hlu *hlunumber*

Specifies the host LUN number, which is the number that the system server will see. *hlunumber* is the virtual LUN number. (Virtual LUN numbers let a host whose operating system has LUN limit or LUN booting restrictions use a number other than the actual system LUN ID (`-alu` below), which might be out of the operating system's range, to access a LUN.) This switch is required along with `-alu` and either the `-gname` or `-uid` switch.

If you have not assigned a host LUN number of 0 and if the `arraycommpath` setting is 1, then after the bus is rescanned, the system will display an entry for LUN 0 even though you have not assigned a host LUN 0. This occurs because the `arraycommpath` setting of 1 creates a virtual LUN 0 for communication with the system.

-alu *alunumber*

Specifies the system LUN number, the actual LUN number internal to the system. This is the number assigned when the LUN was bound. *alunumber* specifies the logical unit number. This switch is required along with `-hlu` and either the `-gname` or `-uid` switch.

`storagegroup` command, `-addsnapshot` switch `-addsnapshot -gname StorageGroupName | -uid StorageGroupUID -hlu hlunumber -snapshotname snapshotname | -snapshotuid snapshotuid`

Adds a snapshot to a storage group. The storage group you choose should be different from the source LUN storage group. You must use three of the following switches with `-addsnapshot`:

-gname *StorageGroupName*

Sets the name of the storage group. *StorageGroupName* specifies a user-defined storage group name; or

-uid *StorageGroupUID*

Specifies the storage group unique identification number.

- For Fibre Channel systems, the UID uses the World Wide Name (WWN) format such as: 10:00:00:00:C9:20:DE:0C:10:00:00:00:C9:20:DE:0C.
- For iSCSI systems, the UID uses the initiator name which uses the IQN format such as: iqn.1992-04.com.emc:cx.apm00034901792.b0.

-hlu *hlunumber*

Specifies the host LUN number for the snapshot, which is the unique number that the system server will see. *hlunumber* is the virtual LUN number.

-snapshotname *snapshotname* | -snapshotuid *snapshotuid*

Specifies the snapshot name or unique ID number (World Wide Name, WWN). One of these switches is required along with the `-gname` or `-uid` switch.

`storagegroup` command, `-chgname` switch `-chgname -gname StorageGroupName | -uid StorageGroupUID -newgname NewName [-o]`

Renames the storage group. You must use two of the following switches with `-chgname`:

-gname *StorageGroupName*

Specifies the name of the storage group. *StorageGroupName* specifies a user-defined storage group name.

-uid *StorageGroupUID*

Specifies the storage group unique identification number.

- For Fibre Channel systems, the UID uses the World Wide Name (WWN) format such as: 10:00:00:00:C9:20:DE:0C:10:00:00:00:C9:20:DE:0C.
- For iSCSI systems, the UID uses the initiator name which uses the IQN format such as: iqn.1992-04.com.emc.cx.apm00034901792.b0.

-newgname *NewName*

Sets the new storage group name for the storage group (specified by the `-gname` or `-uid` switch). *NewName* specifies the new name for the specified storage group.

-o

Executes the command without prompting for confirmation.

`storagegroup command, -connecthost switch -connecthost -host HostName -gname StorageGroupName [-uid StorageGroupUID [-o]]`

Connects the specified host to the specified storage group. It connects all of the host's HBA/SP combinations to the storage group by assigning the host's initiator records to the storage group. To assign a storage group to a specific HBA, use `storagegroup -setpath`.

You may use the `port -list` command to get a list of registered HBAs. If nothing is registered, use the `navicli register` command to register an HBA.

You must specify the host using the `-host` switch, and the storage group using the `-gname` switch or the `-uid` switch. If you include the `-o` switch, the command will be executed without asking for user confirmation. Also, see the example in [unbind on page 181](#).

-host *HostName*

Specifies the user-defined hostname.

-gname *StorageGroupName*

StorageGroupName specifies the storage group name.

-uid *StorageGroupUID*

StorageGroupUID specifies the unique ID.

`storagegroup command, -create switch -create -gname StorageGroupName`

Creates a storage group. You must use the `-gname` switch with `-create`. Also see the example in [unbind on page 181](#).

-gname *StorageGroupName*

StorageGroupName specifies a user-defined storage group name.

`storagegroup command, -destroy switch -destroy -gname StorageGroupName [-uid StorageGroupUID [-o]]`

Destroys the specified storage group. If you include the `-o` switch, the command will be executed without asking for user confirmation. You must use one of the following switches with `-destroy`:

-gname *StorageGroupName*

StorageGroupName specifies a user-defined storage group name.

-uid *StorageGroupUID*

Specifies the storage group UID.

storagegroup command, **-disconnecthost switch** **-disconnecthost -host *HostName* [-o]**

Disconnects all of the specified host's HBA/SP combinations from the currently connected storage group, and reconnects them to ~management. You must specify the host using the **-host** switch. If the **-o** switch is present, the command will be executed without asking for user confirmation.

-host *HostName*

HostName specifies the user-defined hostname.

storagegroup command, **-enable switch** **[-o]**

Enables data access control. Sets the default storage group to ~management. If the **-o** switch is present, the command will be executed without asking for user confirmation.

storagegroup command, **-list switch** **-list [-gname *StorageGroupName*] -uid *StorageGroupUID*] -host**

Displays information about the specified storage group. When no storage group is specified, **-list** displays information about each user-defined storage group in the system. You can use one of the following optional switches with **-list**.

-gname *StorageGroupName*

StorageGroupName specifies a user-defined storage group name.

-host

Displays the hostname of the attached host.

-uid *StorageGroupUID*

Specifies the storage group unique ID.

storagegroup command, **-removehlu switch** **-removehlu -gname *StorageGroupName*] -uid *StorageGroupUID* -hlu *hlnumber***

Removes the specified host LUN number from the specified storage group. You must use one of the following switches:

-gname *StorageGroupName*

StorageGroupName specifies a user-defined storage group name.

-uid *StorageGroupUID*

Specifies the storage group UID.

-hlu *hlnumber*

hlnumber specifies the logical unit number, which is the number that the system server will see. This switch is required along with either the **-gname** or **-uid** switch.

storagegroup command, **-removesnapshot switch** **-removesnapshot -gname *StorageGroupName*] -uid *StorageGroupUID* -snapshotname *snapshotname*] -snapshotwn *snapshotUID***

Removes a snapshot from a storage group. You must use one of the four following switches with **-removesnapshot**:

-gname *StorageGroupName*

StorageGroupName specifies the user-defined storage group name.

-uid *StorageGroupUID*

Specifies the storage group unique ID.

-snapshotname *snapshotname* | -snapshotwwn *snapshotUID*

Specifies the snapshot name or UID number (World Wide Name, WWN). One of these switches is required.

`storagegroup command, -sethost switch -sethost [-host hostname | -ip host-ip-address [-arraycommpath 0|1] [-failovermode 0|1|2 |3|4] [-type initiator-type] [-unitserialnumber array|lun] [-o]`

Creates or changes certain settings for all initiators (HBAs) in the specified host that are connected to a specific system. You must specify the host (`-host` switch), unless you are issuing the command in the current host. To obtain `-sethost` status, use the `port list` command. The `-sethost` function has subfunction switches as follows.

Note

For AX4-5 series systems, the `storagegroup` command with the `-sethost` switch will accept nondefault values for the following switches: `-failovermode`, `-arraycommpath`, `-type`, and `-unitserialnumber`.

-arraycommpath 0|1

`storagegroup -sethost` with `-arraycommpath` modifies one or more communication paths between the server and system. A communication path is a logical path between the host and the system. It may be needed when there is no LUN owned by an SP through which the host recognizes the system and in the storage group through which the host recognizes the system.

Changing the `arraycommpath` setting may force the system to reboot. The CLI will issue a confirmation request that indicates whether a system reboot is required. However, the CLI will not issue a confirmation request if you include the `-o` switch.

⚠ CAUTION

Do not issue this command option unless you know the appropriate value for the type of failover software running on the servers connected to the specified storage group. The command is designed for transition from one type of failover software to another. Changing the `arraycommpath` to the wrong value will make the storage group inaccessible to any connected servers.

This form of `-arraycommpath` works for a host. For an entire system, or unshared storage, use the command `arraycommpath`.

To discover the current setting of `arraycommpath`, use the `port -list -all` or `port -list -arraycommpath` command.

0

Disables the communication path previously enabled.

1

Enables a communication path.

Example with `arraycommpath`:

```
naviseccli -h ssl_spa storagegroup -sethost -hostname server101
-arraycommpath 1
```

-failovermode 0|1|2|3|4

storagegroup -sethost with -failovermode sets the type of trespass needed for failover software. Use this command with certain types of failover software as suggested in the failover software release notice or manual.

Changing the failovermode setting may force the system to reboot. The CLI will issue a confirmation request that indicates whether a system reboot is required. However, the CLI will not issue a confirmation request if you include the -o switch.

CAUTION

Do not issue the failovermode command option unless you know the appropriate value for the type of failover software running on the servers connected to the specified storage group. The command is designed for transition from one type of failover software to another. Changing the failovermode to the wrong value will make the storage group inaccessible to any connected servers.

This method of setting -failovermode works for systems with storage group only. For systems without storage group, use the command failovermode (see [failovermode on page 61](#)).

To discover the current setting of failovermode, use the port -list -failovermode or port -list -all command.

Note

For information about when to use these failover modes, refer to the *Installation Roadmap for CLARiiON Storage Systems*, which is on EMC Online Support.

0

Failover mode 0.

1

Failover mode 1.

2

Failover mode 2.

3

Failover mode 3.

4

Failover mode 4.

Example with failovermode:

```
naviseccli -h ssl_spa storagegroup -sethost -failovermode 1
-hostname server102
```

-hostname [host]

Specifies the HBA owner's hostname. You must specify the host's hostname or its IP address (-ip host-ip-address switch).

-ip *host-ip-address*

Specifies the IP address of the host, in the form *xx.xx.xx.xx*. You must specify either the IP address or the hostname. The IP address can be specified as an IPv4 or IPv6 address. If an incorrect IP address is specified, the error message is displayed as follows:

```
The IP address must be specified as an IPv4 address
in dotted-decimal format or as an IPv6 address in
colon-hexadecimal notation.
```

-type *initiator*

Specifies an initiator type to override the system default. The CLI uses the default, 3, for the target if omitted. Type values (specify 0x to indicate hex) are as follows:

0x2

HP (auto trespass)

0x3

Open (default).

0x9

SGI

0xa

HP (no auto trespass)

0x13

Dell

0x16

Fujitsu

0x1c

Compaq Tru64

-unitserialnumber [array|lun]

The `unitserialnumber` command displays or sets the VDPP80 behavior mode, which defines how subsequent VPDP80 requests will be populated with the system serial number. You might need to do this for software such as SunCluster 3. You can discover the mode with the `port -list -unitserialnumber` or `port -list -all` command. To set this mode for a host, see `storagegroup -sethost` above.

array

Sets the behavior mode to system.

lun

Sets the behavior mode to LUN.

Example with `unitserialnumber`:

```
naviseccli -h ssl_spa storagegroup -sethost -unitserialnumber
lun
```

```
storagegroup command, -setpath switch -setpath [-gname
StorageGroupName| -uid StorageGroupUID] [-arraycommpath 0|1] [-
```

```
failovermode 0|1|2 |3|4] [-hbauid hbauid] [-host hostname] [-ip ip-address]
[-sp a|b] [-spport 0|1] [-unitserialnumber array|lun] [-o]
```

Creates or changes settings for a specific initiator (HBA). You must specify the HBA (-hba switch), SP (-sp), and SP port (-spport). You can identify the storage group by name (-gname switch) or unique ID (-uid switch). If you omit both switches, the CLI uses the default path. To obtain -setpath status, use the port list command. The -setpath function has many function switches as follows.

-arraycommpath 0|1

storagegroup -setpath with -arraycommpath modifies one or more communication paths between the server and system. It may be needed when there is no LUN owned by an SP through which the host recognizes the system and in the storage group through which the host recognizes the system. Changing the arraypath setting forces the system to reboot.

This form of -arraycommpath works for systems with storage group only. For systems without storage group, use the command arraycommpath (see [arraycommpath on page 28](#)).

To discover the current setting of arraycommpath, use the port -list -arraycommpath or port -list -all command.

0

Disables the communication path previously enabled.

1

Enables a communication path.

Example with arraycommpath:

```
naviseccli -h ss1_spa storagegroup -setpath -gname Group1
-hbauid 10:00:00:00:C9:20:DE:0C:10:00:00:00:C9:20:DE:0C
-sp a -spport 0 -arraycommpath 1
```

-failovermode 0|1|2 |3|4

storagegroup -setpath with -failovermode sets the type of trespass needed for failover software. Use this command with certain types of failover software as suggested in the *Installation Roadmap for CLARiiON Storage Systems*, which is on EMC Online Support.

Changing the failovermode setting may force the system to reboot. The CLI will issue a confirmation request that indicates whether a system reboot is required. However, the CLI will not issue a confirmation request if you specified the -o switch.

This form of -failovermode works for one storage group. To set the entire system mode, use the command failovermode.

To discover the current setting of failovermode, use the port -list -failovermode or port -list -all command.

0

Failover mode 0.

1

Failover mode 1.

- 2
Failover mode 2.
- 3
Failover mode 3.
- 4
Failover mode 4.
- o
Executes the command without prompting for confirmation.

Example with failovermode:

```
naviseccli -h ssl_spa storagegroup -setpath -gname
Group1 -hbauid
10:00:00:00:C9:20:DE:0C:10:00:00:00:C9:20:DE:0C
-sp a -sport 0 -failovermode 1
```

-gname *StorageGroupName*

StorageGroupName specifies a user-defined storage group name. Use this or the storage group unique ID (below).

-hbauid *hbauid*

Specifies the HBA unique ID (WWN). Required.

-host [*hostname*]

Specifies the HBA owner's hostname. The system uses the hostname (or `-ip address` switch) for any subsequent `storagegroup -connecthost` and `-disconnecthost` commands you issue. If you omit this, the system uses the HBA UID. Not required.

-ip *address*

Specifies the HBA owner's IP address (server's IP address). You can specify the IP address as the IPv4 or IPv6 address. Not required. See `-hostname` switch.

-sp *a|b*

Specifies the SP that owns the storage group. Required.

-sport *0|1*

Specifies the SP port. Required.

-uid *StorageGroupUID*

Specifies the storage group unique ID.

-unitserialnumber *array|lun*

Sets the unit serial number operation mode. Use this switch if the software instructs you to.

-o

Executes the command without prompting for confirmation.

`storagegroup command, -shareable switch -shareable [-gname StorageGroupName [yes|no] | -uid StorageGroupUID [yes|no]]`

Specifies whether more than one host can share the storage group. You must use one of the following switches with `-shareable`, followed by a yes or a no. A yes or no indicates whether the storage group can be shared. Only Unisphere can enforce sharing.

-gname *StorageGroupName* yes|no

Sets the storage group named *StorageGroupName* to be shareable (yes) or not shareable (no).

-uid *StorageGroupUID* yes|no

Sets the storage group with unique ID *StorageGroupUID* to be shareable (yes) or not shareable (no).

EXAMPLE #1

In the following example, Unisphere connects all of `ss1_spa`'s HBA-SP combinations to the storage group mail.

```
naviseccli -h ss1_spa storagegroup -gname mail -connecthost -o
```

EXAMPLE #2

In the following example, the `-list` command displays storage group information.

```
naviseccli -h ss1_spa storagegroup -list
```

```
Storage Group Name: Blade 150
Storage Group UID:
34:C7:72:A9:2D:DB:D8:11:AF:5E:08:00:1B:43:73:EE
HBA/SP Pairs:

HBA UID                               SP Name
SPPort
-----                               -
-----
20:00:00:00:C9:39:40:C7:10:00:00:00:C9:39:40:C7 SP B
0
20:00:00:00:C9:39:40:C8:10:00:00:00:C9:39:40:C8 SP A
0

HLU/ALU Pairs:

HLU Number  ALU Number
-----
0             15
1             14
2             13
3             11
4             10
5             9
6             7
7             6
8             5
9             3
10            2
11            1
Shareable: YES
```

EXAMPLE #3

The following examples creates two storage groups, connects them to separate hosts, and adds a LUN to each one.

```
naviseccli -h ss_spa storagegroup -create -gname Production
```

```
naviseccli -h ss_spa storagegroup -create -gname Backup
```

```
naviseccli -h ss_spa storagegroup -connecthost -o -host ServerABC
-gname Production
```

```
naviseccli -h ss_spa storagegroup -connecthost -o -host ServerXYZ
-gname Backup
```

```
naviseccli -h ss_spa storagegroup -addhlu -gname Production -hlu
20 -alu 20
```

```
naviseccli -h ss_spa storagegroup -addhlu -gname Backup -hlu 21
-alu 21
```

EXAMPLE #4

The following example of `storagegroup setpath` and `storagegroup sethost` with `failovermode` sets the type of trespass needed for failover software. Also, the failover mode option is set to 4 in `sethost` and `setpath` switches.

```
naviseccli -h 10. 4. 78. 184 -user a -password a -scope 0 storagegroup
-sethost -ip 10.4.78.184 -failovermode 4
```

```
naviseccli -h 10. 4. 78. 184 -user a -password a -scope 0 storagegroup
-setpath -hbauid xxxxx-sp a -spport xxxxx -failovermode 4
```

EXAMPLE #5

The following example specifies the form of `failovermode` for one storage group. For the options `-failovermode` and `-all`, the command displays the new failover mode as 4 under the HBA information section when the failover mode is active/active.

```
naviseccli -h 10. 4. 78. 184 -user a -password a -scope 0 port -list
-failovermode
```

```
Information about each port of this HBA
SP Name : SP A
SP Port ID : 0
Failover Mode : 4
```

```
naviseccli -h 10. 4. 78. 186 -user a -password a -scope 0 port -list
-all
```

```
HBA UID: 20 : 20 : 20 : 20 : 20 : 20 : 20 : 20 : 20 : 20
: 20 : 20 : 20 : 20 : 20 : 20
Server Name : TCS_Dell
Server IP Address : 10 . 4 . 78 . 159
HBA Model Description: 2
HBA Vendor Description : DELL
HBA Device Driver Name : N/A
Information about each port of this HBA :
SP Name : SP A
SP Port ID : 0
HBA Devicename : N/A
Trusted : NO
Logged In : NO
Defined : YES
Initiator Type : 3
StorageGroup Name : None
Array Commpath : 1
```

```
Failover mode : 4
Unit serial number : Array
```

systemtype

Sets the system type.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli systemtype` command changes the configured system type. Write caching must be disabled for this command to work. Changing the system type causes an SP to restart.

If you issue the command without arguments, the CLI will display both the current and configured system types.

Note

Do not change the system type of any system running storage group. Do not change the system type of a system unless it has been quiesced. The agent must be running and managing the system.

Depending on the type of HP-UX operating system, you may want auto-trespass to be on or off.

SYNTAX

```
systemtype -config [type] [-o]
```

OPTIONS

`-config [type]`

Sets the system type to one of the following:

A

HP-UX type with auto-trespass off; use this if your HP-UX system does not have the PVLINKS feature.

2

HP-UX type with auto-trespass on; use this if your HP-UX system has the PVLINKS feature.

3

Open system (default)

9

SGI (CX600)

13

Dell

16

Fujitsu

1C

HP type

`-o`

Executes the command without prompting for confirmation.

EXAMPLE #1

```
naviseccli -h ss1_spa systemtype
```

```
The configured system type is: 0x3
The current system type is: 0x3
```

EXAMPLE #2

```
naviseccli -h ss1_spa systemtype -config 2
```

```
In order for this command to execute, write cache
must be disabled and memory size set to zero:
```

trespass

Trespasses one or more LUNs to the control of an SP.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command trespasses one or more LUNs to the control of an SP.

Note

Not applicable to single-SP systems.

The command attempts to trespass all specified LUNs to the SP you specify with the `-h` switch.

Note

If trespass is executed while I/O processes are active, I/O errors may occur.

SYNTAX

```
trespass [all | lun lun-number | mine]
```

OPTIONS

`all`

Trespasses all LUNs in the system.

`lun lun-number`

Trespasses one LUN (specified by *lun-number*).

`mine`

Performs whatever trespassing is necessary to take all of the LUNs that have this SP as a default owner.

EXAMPLE

```
naviseccli -h ss1_spa trespass all
```

OUTPUT

There is no output. Errors are printed to `stderr`.

unbind

Deconfigures a LUN.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli unbind` command deconfigures LUNs from their current logical configuration. The `unbind` command destroys all data on the LUN; therefore, when you execute an `unbind` command, you are prompted to verify that you really want to unbind the LUN(s). If you do not want the prompt to appear, you can override prompting using the `-o` switch.

The command can unbind only LUNs owned by the SP you specify with `-h`.

You can unbind multiple LUNs at the same time.

Note

You cannot use this command to unbind thin LUNs.

SYNTAX

```
unbind lun [-o]
```

OPTIONS

lun

Specifies the LUN to be deconfigured.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE #1

This command destroys (deconfigures) LUN 3 without prompting you for permission, and frees its disks to be reconfigured.

```
naviseccli -h ss1_spa unbind 3 -o
```

EXAMPLE #2

This command destroys (deconfigures) LUNs 0, 1, 2, 3, and 4 without prompting you for permission, and frees its disks to be reconfigured.

```
naviseccli -h ss1_spa unbind 0 1 2 3 4 -o
```

unitserialnumber

Displays or changes the VDPP80 (port 80) behavior mode.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli unitserialnumber` command displays or sets the VDPP80 behavior mode, which defines how subsequent VPDP80 requests will be populated with the system serial number. You might need to do this for software such as SunCluster 3.

To set this mode for a specific storage group, see the `storagegroup` command in `storagegroup`.

SYNTAX

```
unitserialnumber [array|lun] [-o]
```

OPTIONS

array

Sets the behavior mode to system.

lun

Sets the behavior mode to LUN.

-o

Executes the command without prompting for confirmation.

EXAMPLE

These two commands display the current behavior mode and then set the mode to system mode.

```
naviseccli -h ssl_spa unitserialnumber
```

```
Current unit serial number mode is: LUN
```

```
naviseccli -h ssl_spa unitserialnumber array
```

LUN Expansion (metaLUN) commands

This section explains the VNX CLI for Block commands that manage Unisphere metaLUNs.

MetaLUNs overview

A metaLUN is a type of LUN whose maximum capacity is the combined capacities of all the LUNs that compose it. The metaLUN feature lets you dynamically expand the capacity of a single LUN (primary LUN) into a larger unit called a metaLUN. You do this by adding LUNs to the primary LUN. You can also add LUNs to a metaLUN to further increase its capacity. Like a LUN, a metaLUN can belong to a storage group, and can participate in SnapView, MirrorView and SAN Copy sessions.


During the expansion process, the host has access to any existing data on the LUN or metaLUN, but it does not have access to any added capacity until the expansion is complete. Whether you can actually use the increased user capacity of the metaLUN depends on the operating system running on the servers connected to the system.

A metaLUN can include multiple sets of LUNs that are striped together and are independent of other LUNs in the metaLUN. Each set of striped LUNs is called a component. All metaLUNs contain at least one component that includes the primary LUN and one or more LUNs. Any data that gets written to a metaLUN component is striped across all the LUNs in the component.

You can expand a LUN or metaLUN in two ways — stripe expansion or concatenate expansion. A stripe expansion takes the existing data on the LUN or metaLUN, and restripes (redistributes) it across the existing LUNs and the new LUNs you are adding. The stripe expansion may take a long time to complete. A concatenate expansion creates a new metaLUN component that includes the new LUNs and appends this component to the end of the existing LUN or metaLUN. There is no restriping of data between the original storage and the new LUNs. The concatenate operation completes immediately.


[Sample stripe expansion on page 183](#) depicts a LUN using stripe expansion to add two LUNs. The result is a metaLUN with a single component consisting of all three LUNs with the data on the primary LUN re-striped across all three LUNs.

Figure 1 Sample stripe expansion



[Sample concatenate expansion on page 183](#) depicts a LUN using concatenate expansion to add two LUNs. The result is a metaLUN with two components. The first component is the primary LUN. The second component consists of two new LUNs striped together, but not striped to the first component.

Figure 2 Sample concatenate expansion



Note

EMC does not recommend that you expand LUN capacity by concatenating LUNs of different RAID types. Do this only in an emergency situation when you need to add capacity to a LUN and you do not have LUNs of the same RAID type or the disk capacity to bind new ones. Concatenating metaLUN components with a variety of RAID types could impact the performance of the resulting metaLUN. Once you expand a LUN, you cannot change the RAID type of any of its components without destroying the metaLUN. Destroying a metaLUN unbinds all LUNs in the metaLUN, and therefore causes data to be lost.

CLI metaLUN command functions

The `naviseccli metaLUN` command functions, listed alphabetically, are:

- `metalun -destroy`
- `metalun -expand`
- `metalun -info`
- `metalun -list`
- `metalun -modify`

The CLI metaLUN functions are listed functionally, in the order you would perform them, in [naviseccli metaLUN command functions, listed in order of use on page 183](#).

Table 10 `naviseccli metaLUN` command functions, listed in order of use

Essential functions (in order performed)	
<code>metalun -expand</code>	Creates a metaLUN or expands one
<code>metalun -destroy</code>	Destroys a metaLUN and then unbinds all LUNs that were part of the metaLUN

Table 10 naviseccli metaLUN command functions, listed in order of use (continued)

Optional and status functions (alphabetically)	
<code>metalun -info</code>	Provides metaLUN information with respect to a specific system
<code>metalun -list</code>	Lists the properties of existing metaLUNs
Optional reconfiguration functions (alphabetically)	
<code>metalun -modify</code>	Modifies certain metaLUN properties

metalun -destroy

Destroys a metaLUN and unbinds all LUNs that were part of the metaLUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli metalun` command with the `-destroy` switch destroys an existing metaLUN and then unbinds all LUNs that were part of the metaLUN. You will lose all data on the LUNs.

Note

You cannot destroy a metaLUN that belongs to a storage group. You must remove the metaLUN from the storage group in order to destroy it. You cannot destroy any metaLUNs that are currently participating in a snapshot, clone, mirror or SAN Copy session.

SYNTAX

```
metalun -destroy -metalun number|WWN [-o]
```

OPTIONS

`-metalun number| WWN`

Specifies the unique number or World Wide Name of the metaLUN you want to destroy. Use only one of the identifiers with the `-metalun` command. Use the `metalun -list` command to determine the number or WWN of the metaLUN (see [metalun -list on page 191](#)).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command destroys the metaLUN with the number 20.

```
naviseccli metalun -destroy 20
```

metalun -expand

Creates a new metaLUN or expands the capacity of an existing one

PREREQUISITES

This section describes the general requirements and recommendations for both a stripe and concatenate expansion, as well as those specific to one or the other.

General expansion requirements and recommendations

- Each LUN within a metaLUN must be either a redundant (RAID 6, RAID 5, RAID 3, RAID 1, RAID 1/0) or non-redundant (RAID 0, disk) RAID type.

- All disks within a metaLUN must be the same type.
- You can expand only the last component of a metaLUN.
- Expansion LUNs cannot be participating in a SnapView, MirrorView or SAN Copy session.
- The `metalun -expand` command will prohibit metaLUN creation and expansion if the element size of the selected LUN(s) differ from the primary LUN. The system will display the following error message:

```
The selected LUN(s) element size must be the same as the LUN
being expanded. Please select LUN(s) that match the primary
LUN's element size.
```

Stripe expansion requirements and recommendations

All LUNs within a striped metaLUN or metaLUN component must be the same size and RAID type, and we recommend that each LUN:

- be from a different RAID group.
- includes the same number of disks.
- has the same stripe element size.

Concatenate expansion requirements and recommendations

All LUNs within a concatenated metaLUN component must be the same size and RAID type, but can differ from the size and RAID type of LUNs within other components.

Note

EMC does not recommend that you expand LUN capacity by concatenating LUNs of different RAID types. Do this only in an emergency situation when you need to add capacity to a LUN and you do not have LUNs of the same RAID type or the disk capacity to bind new ones. Concatenating metaLUN components with a variety of RAID types could impact the performance of the resulting metaLUN. Once you expand a LUN, you cannot change the RAID type of any of its components without destroying the metaLUN. Destroying a metaLUN unbinds all LUNs in the metaLUN, and therefore causes data to be lost.

DESCRIPTION

The `metalun -expand` command is used with `naviseccli` (see [naviseccli on page 20](#)). The `naviseccli metalun` command with the `-expand` function creates a new metaLUN or expands the capacity of an existing one. There are two distinct versions of this command, one for Unisphere and one for Navisphere Express. When using this command with Navisphere Express, the syntax is simpler and does not require binding the destination LUN first. The process is different with Unisphere and requires you to bind the destination LUN before running the command. See the Prerequisite section above for requirements and recommendations.

Note

Thin LUNs do not support metaLUNs, so you cannot create metaLUNs using thin LUNs. For raw device mapping volumes (RDMs), if you are running VMware ESX Server, after you create or expand the metaLUN, you must rescan the bus at the ESX Server level. You can also rescan at the virtual machine level (optional). You then must expand the file system at the virtual machine level. ESX Server and the virtual machine will now recognize the new or expanded metaLUN.

For configuration specifications, go to the E-Lab™ Interoperability Navigator on EMC Online Support.

SYNTAX

```
metalun -expand -base number|WWN -lus number|WWN [-autoassign
0|1] [-bindoffset sector-number] [-cap capacity]
[-defaultowner A|B] [-elszm elementSizeMultiplier]
[-expansionrate low|medium|high|ASAP|value]
[-name metaLUNname] [-preservedata 1|0] [-nogrow] [-o]
[-sq tb|gb|mb|bc] [-type C|S]
```

```
metalun -expand -base number|WWN -lus number|WWN [-type C|S]
[-sq tb|gb|mb|bc] [-cap number] [-nogrow] [-o]
```

```
metalun -expand -base number|WWN [-sq tb|gb|mb -cap number]
[-percent number] [-o]
```

OPTIONS

-base *number|WWN*

Specifies the unique number or World Wide Name of the LUN or metaLUN you want to expand. Use only one of the identifiers with the `-base` switch. Use the `metalun -list` command to determine the number or WWN of the LUN or metaLUN (see [metalun -list on page 191](#)).

When you specify a LUN with the `-base` command, the CLI creates a new metaLUN with the number and WWN of the primary LUN. Once the CLI creates the new metaLUN, it assigns a new WWN to the original primary LUN and makes it a private LUN. When expanding a single LUN, you can use any or all of the `-expand` switches.

Note

The primary LUN can belong to a storage group. Expansion LUNs cannot belong to a storage group.

When you specify an existing metaLUN with the `-base` command, the CLI expands the capacity of the existing metaLUN. The following switches are not valid during the expansion of an existing metaLUN: `-name`, `-preservedata`, `-expansionrate`, `-defaultowner`, `-autoassign`, `-elsz`, and `-bindoffset`. The expanded metaLUN retains the values of the original metaLUN. After the metaLUN is expanded, use the `metalun -modify` command to change the metaLUN using `-name`, `-expansionrate`, `-defaultowner`, or `-autoassign`. The `-preservedata`, `-elsz` and `-bindoffset` values cannot be changed.

-lus *lunnumber|WWN*

Specifies the LUN number or World Wide Name of the LUNs you want to add to the LUN or metaLUN (expansion LUNs). Use only one of the identifiers per LUN with the `-`

`lus` command. Use the `metalun -list` command to determine the number or WWN of the LUNs (see [metalun -list on page 191](#)).

Note

Expansion LUNs cannot be participating in a SnapView (snapshot or clones), MirrorView, or SAN Copy session.

`-autoassign 0|1`

Enables or disables autoassign for a metaLUN. Specify 1 to enable autoassign and 0 to disable it. If you do not specify an autoassign value, it defaults to the value of the primary LUN.

Autoassign controls the ownership of the LUN when an SP fails in a system with two SPs. With autoassign enabled, if the SP that owns the LUN fails and the server tries to access that LUN through the second SP, the second SP assumes ownership of the LUN to enable access. The second SP continues to own the LUN until the failed SP is replaced and the system is powered up. Then, ownership of the LUN returns to its default owner. If autoassign is disabled in the previous situation, the second SP does not assume ownership of the LUN, and access to the LUN does not occur.

If you are running failover software on a Windows or UNIX server connected to the system, you must disable autoassign for all LUNs that you want to fail over to the working SP when an SP fails.

`-bindoffset sector-number`

Note

Available only if the `preservedata` switch is set to 0.

EMC does not recommend using a metaLUN bindoffset value if the primary LUN already has one.

Automatically aligns the logical sector number of the metaLUN to a stripe boundary on the primary LUN. The value to enter is the number of sectors by which user data is offset due to file system or partition metadata. Correct application of the bind offset should result in improved performance. Use a decimal, not hexadecimal, number for sector-number. If you do not specify a bind offset value, it defaults to 0.

`-cap capacity`

Specifies the capacity of usable space in the metaLUN. The capacity must be smaller than the maximum capacity of the metaLUN components, and greater than or equal to the current capacity (user capacity of the LUN or metaLUN before adding LUNs). The default value is the maximum capacity of the metaLUN (total capacity of the existing LUNs and any new LUNs).

Note

If you want to expand a LUN or metaLUN that is participating in a SnapView (snapshot or clone), MirrorView, or SAN Copy session, you must use the `-nogrow` switch with the `-expand -cap` command. This command expands the capacity of the LUN or metaLUN to the specified value, but keeps the available user capacity of the base (LUN or metaLUN) the same as it was before adding LUNs. In order for the host to have access to the added capacity, you must wait for the session to complete and then follow the instructions in [Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN on page 197](#).

`-defaultowner A|B`

Specifies the SP that will assume ownership of the metaLUN when the system is powered up. If you do not specify a default owner, it defaults to the owner of the primary LUN.

`-elszm elementSizeMultiplier` (Not supported on AX4-5 series systems running Unisphere)

Sets the stripe element size for a metaLUN component. You determine this by multiplying the stripe size of the first LUN in each component by the element size multiplier. The stripe size of the first LUN = stripe element size of the LUN multiplied by the number of disks in the LUN. Valid multiplier values are 1 through 255. The default element size multiplier is 4.

`-expansionrate low|medium|high|ASAP|value`

Specifies the relative importance of making additional LUN capacity available to the host. It determines the amount of resources the SP devotes to re-striping the existing data on the original LUNs across the original LUNs and the new LUNs instead of to normal I/O activity. This rate is available only when you first create a metaLUN and is important only for a stripe expansion. Valid rebuild priorities are ASAP (default) or 3, High or 2, Medium or 1, and Low, or 0.

An expansion rate of ASAP or High re-stripes the data across the metaLUN faster than one with Medium or Low priority, but may degrade system performance.

`-name metaLUN-name`

Specifies the name of the metaLUN. If you do not specify a name, it defaults to the name of the primary LUN.

`-nogrow`

Keeps the available user capacity of the metaLUN the same as it was before adding LUNs. You cannot use the added capacity until you change the capacity using the `-modify` switch.

Note

For a stripe expansion, any existing data is restriped across the original storage and any new LUNs, but you cannot access the added capacity at this time.

`-o`

Executes the command without prompting for confirmation.

`-preservedata 1|0`

Use with the `-type S` command (see below). Allows you to preserve or overwrite any existing data on the primary LUN, only if the primary LUN is outside a storage group.

Note

You cannot use this command with the `-type C` switch.

If the primary LUN is outside of a storage group, do one of the following:

- Specify 1 to enable stripe expansion and preserve the data on the primary LUN. Note that this operation could take a while. The default value is 1.

Note

You must always specify 1 when the primary LUN is in a storage group. The CLI returns an error message if you specify 0 in this case.

- Specify 0 to enable stripe expansion and overwrite the existing data on the primary LUN. Note that you will lose the original data on the LUN.

`-sq tb|gb|mb|bc`

Specifies the size qualifier for the capacity of the metaLUN component.

where:

tb equals terabytes, gb equals gigabytes, mb equals megabytes and bc equals block count.

Note

The default qualifier is block count.

-type C|S

Sets the expansion type for the metaLUN to concatenate or stripe. The default expansion type is stripe.

C sets the expansion type to concatenate. A concatenate expansion appends the new LUNs to the end of the primary LUN component or metaLUN as a single, separate striped component. There is no striping between the original storage and the new LUNs.

S sets the expansion type to stripe. For a list of requirements for stripe expansion, see the Prerequisite section above.

If you are expanding a LUN (the primary LUN), the `-type S` switch creates a metaLUN with one component that includes the primary LUN and any LUNs that you add. Any data on the primary LUN is re-striped across the primary LUN and the new LUNs.

Note

If the primary LUN has existing data, and it is outside of a storage group, refer to the `-preservedata` switch (see above) for additional options.

You can expand only the last component of an existing metaLUN. In this case, the `-type S` switch re-stripes any existing data across the original LUNs in the component and any new LUNs that you add. The new LUNs become part of that metaLUN component.

-base *number*|*WWN*

Mandatory switch, supported with Navisphere Express, that specifies the source LUN or metaLUN you want to expand. You can specify the base number or World Wide Name of the LUN.

-cap *number*

Mandatory switch, supported with Navisphere Express, that specifies the additional storage capacity you can add to the LUN or metaLUN. The size qualifier for the `-cap number` is the `-sq` switch.

-sq *tb*|*gb*|*mb*

Mandatory switch, supported with Navisphere Express, that specifies the size qualifier for the capacity of the metaLUN component.

where:

tb equals terabytes, gb equals gigabytes, and mb equals megabytes.

For example, if you are expanding a 20 GB LUN, and specify `-sq gb -cap 20` in the command, the resulting metaLUN is 40 GB.

-percent *number*

Optional switch for Navisphere Express that specifies the amount of additional storage you want to expand by, as a percentage of the current size of the LUN. The *number* should always be greater than 0. This switch can be used instead of `-sq` and `-cap number` switches.

For example, if you are expanding a 20 GB LUN, and specify `-percent 50` in the command, the resulting metaLUN is 30 GB.

EXAMPLE

For Unisphere Manager, this command creates a metaLUN with the name `newMetaLUN`, which includes primary LUN 20 and LUN 42. Since the default expansion type is stripe, this metaLUN is striped.

```
naviseccli metalun -expand -base 20 -lus 42 -name newMetaLUN
```

metalun -info

Provides general information about a system's metaLUN environment.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli metalun` command with the `-info` function and no other options displays all general information about the metaLUN environment for a system. Issuing this command with one or more options displays the information based on the options issued.

SYNTAX

```
metalun -info [-canexpand] [-nummetaluns] [-maxmetaluns]
[-maxcomponents] [-maxluns] [-availableluns]
[-lunsinmeta]
```

OPTIONS

- `-canexpand`
Specifies whether you can or cannot expand any LUNs or metaLUNs on the system. Valid values are Yes or No.
- `-nummetaluns`
Displays the number of metaLUNs in the system.
- `-maxmetaluns`
Displays the maximum number of metaLUNs that the system will support (based on system type). See the *EMC Unisphere Release Notes* for currently supported values.
- `-maxcomponents`
Displays the maximum number of components allowed per metaLUN for this system (based on system type). See the *EMC Unisphere Release Notes* for currently supported values.
- `-maxluns`
Displays the maximum number of LUNs per component that this system supports (based on system type). See the *EMC Unisphere Release Notes* for currently supported values.
- `-availableluns`
Lists all LUNs (by LUN number) in the system that are available for expansion.
- `-lunsinmeta`
Lists any LUNs (by LUN number) in the system that are part of a metaLUN.

EXAMPLE

This command requests information about the metaLUN environment for the system with SP 10.14.20.57.

```
naviseccli -address 10.14.20.57 metalun -info
```

OUTPUT

```

Can a LUN or a Meta-LUN be expanded on this system: Yes
Number of Meta-LUNs in the system: 7
Maximum Number of Meta-LUNs per system: 1024
Luns that are participating in Meta-LUNs: 513 512 515 514 517 516
519 518 521 52
0 523 522 525 524
Luns that are available for expansion: 32 43 44 47 48 49 50 51 52 53
54 55 56 5
7 58 59 0
Maximum Number of Components per Meta-LUN: 16
Maximum Number of LUNs per Component: 32

```

The CLI returns status or error information if the command fails.

metalun -list

Lists the properties of existing metaLUNs.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli metalun` command with the `-list` function and no other options lists all property information for all metaLUNs on a system. Issuing this command with one or more options displays the information based on the options issued.

Note

In order to return values for performance data, you must enable statistics logging for each SP in the system. Use VNX CLI for Block or Unisphere. Performance data switches are marked with an asterisk (*).

If you request performance data for a specific metaLUN, make sure that you enter the IP address for the SP that is the current owner of the metaLUN.

SYNTAX

```

metalun -list [-name metaLUN-name] [-metalunwwn metaLUN-wwn]
[-metalunnumber metaLUN-number]
[-all] [-aac*] [-actualusercap] [-autoassign] [-bindoffset]
[-brw*]
[-components*] [-currentowner] [-canexpand] [-defaultowner]
[-drivetype] [-expansionrate] [-elszm] [-isredundant] [-nzcra]
[-percentexp] [-rhist*] [-rwr*] [-sor] [-state] [-totalusercap]
[-whist*]

```

OPTIONS

`-name metaLUN-name`

Lists all information for that specific metaLUN.

Issuing the `-name` command without a metaLUN-name displays the names of all metaLUNs on the system.

`-metalunwwn metaLUN-wwn`

Lists all information for that specific metaLUN.

Issuing the `-metalunwwn` command without a metaLUN-wwn displays the WWNs for all metaLUNs on the system.

`-metalunnumber metaLUN-number`

Lists all information for that specific metaLUN.

Issuing the `-metalunnumber` command without a metaLUN-number displays the numbers for all metaLUNs on the system.

`-aac*`

Displays the active/active counters associated with the metaLUN.

`-actualusercap`

Displays the current user capacity of all metaLUNs on the system.

Issuing the `-actualusercap` command with a metaLUN name, number or WWN displays the user capacity for that specific metaLUN.

`-autoassign`

Displays the current state of the autoassign feature for all metaLUNs on the system (Yes or No).

Issuing the `-autoassign` command with a metaLUN name, number or WWN displays the current state of the autoassign feature for that specific metaLUN.

`-all`

Displays all information for the `metalun -list` command.

`-bindoffset`

Aligns the logical sector number of all metaLUNs on the system to a stripe boundary on the primary LUN.

Issuing the `-bindoffset` command with a metaLUN name, number or WWN returns the value of the bindoffset used for that metaLUN.

`-brw*`

Displays the number of host block read and write requests to all metaLUNs on the system.

Issuing the `-brw` command with a metaLUN name, number or WWN displays the number of host block read and write requests to that specific metaLUN.

`-components*`

Displays the contents of all components in all metaLUN on the system.

Issuing the `-components` command with a metaLUN name, number or WWN displays the contents of all components for that specific metaLUN.

`-currentowner`

Lists the current SP owner of all metaLUNs in the system. If neither SP owns the metaLUN a value of Unknown is returned.

Issuing the `-components` command with a metaLUN name, number or WWN lists the current SP owner of the specified metaLUN.

`-canexpand Yes | No`

Yes

The specified metaLUN can be expanded.

No

The specified metaLUN cannot be expanded.

`-defaultowner`

Lists the default SP owner of all metaLUNs on the system.

Issuing the `-defaultowner` command with a metaLUN name, number or WWN lists the default SP owner of the specified metaLUN.

`-drivetype`

Displays the drive type for all metaLUNs on the system.

Issuing the `-drivetype` command with a metaLUN name, number or WWN displays the drive type for the specified metaLUN.

`-expansionrate`

Displays the expansion rate for all metaLUNs on the system (see the `-expansionrate low|medium|high|ASAP|value` option in [metalun -expand on page 184](#)). Valid values are 0 (Low), 1 (Medium), 2 (High), 3 (ASAP), or N/A if the metaLUN is not currently expanding.

Issuing the `-expansionrate` command with a metaLUN name, number or WWN displays the expansion rate for the specified metaLUN (see the `-expansionrate low|medium|high|ASAP|value` option in [metalun -expand on page 184](#)).

`-elszm`

Displays the element size multiplier for all metaLUNs on the system (see `-elszm elementSizeMultiplier` in [metalun -expand on page 184](#)).

Issuing the `-elszm` command with a metaLUN name, number or WWN displays the element size multiplier for the specified metaLUN (see `-elszm elementSizeMultiplier` in [metalun -expand on page 184](#)).

`-isredundant`

Displays the redundancy of all metaLUNs on the system:

Yes

if it is redundant

No

if it is not redundant

Issuing the `-isredundant` command with a metaLUN name, number or WWN displays the redundancy of the specified metaLUN.

`-nzcra`

Returns the nonzero request count arrivals in SP A and SP B.

`-percentexp`

Displays the percentage complete of the stripe expansion process of all primary LUNs or metaLUNs on the system.

Issuing the `-percentexp` command with a metaLUN name, number or WWN displays the percentage complete of the stripe expansion process of the specified primary LUN or metaLUN.

`-rhist*`

Returns read histogram information for all metaLUNs on the system.

Note

A read histogram is an system of 10 locations that contain the number of reads. Element n of the system contains the number of reads that were larger than or equal to $2n-1$ and less than $2n$ blocks in size. For example, the first bucket contains the number of 1 block reads, the second contains the number of 2-3 block reads, the third contains the number of 4-7 block reads, and so on. The `-rhist` switch also displays the number of read histogram overflows, which are the number of reads that were larger than 512 blocks.

Issuing the `-rhist` command with a metaLUN name, number or WWN returns the read histogram for the specified metaLUN.

`-rwr*`

Displays the number of host read and write requests to all metaLUNs on the system.

Issuing the `-rwr` command with a metaLUN name, number or WWN displays the number of host read and write requests to the specified metaLUN.

`-sor`

Returns the sum of outstanding requests in SP A and SP B.

`-state`

Displays the current state of all metaLUNs on the system.

When issuing the `-state` command with a metaLUN name, number or WWN displays the state of the specified metaLUN. The following are valid metaLUN states:

Normal

The metaLUN is operating normally.

Degraded

Some part of the metaLUN is faulted, but the metaLUN is still able to process I/O.

Shut down

The metaLUN is faulted and unusable.

Expanding

The metaLUN is in the process of expanding (re-striping data across all component LUNs).

`-totalusercap`

Displays the total user capacity of all metaLUNs on the system.

Issuing the `-totalusercap` command with a metaLUN name, number or WWN displays the total user capacity of the specified metaLUN.

`-whist*`

Returns write histogram information for all metaLUNs on the system.

Note

A write histogram is an system of 10 locations that contain the number of writes. Element n of the system contains the number of writes that were larger than or equal to $2n-1$ and less than $2n$ blocks in size. For example, the first bucket contains the number of 1 block writes, the second contains the number of 2-3 block writes, the third contains the number of 4-7 block writes, and so on. The `-whist` switch also displays the number of write histogram overflows, which are the number of writes that were larger than 512 blocks.

Issuing the `-whist` command with a metaLUN name, number or WWN returns the write histogram for the specified metaLUN.

EXAMPLE

This command requests information on all metaLUNs that are currently owned by the SP with IP address 10.14.20.57.

```
naviseccli -address 10.14.20.57 metalun -list
```

OUTPUT

```
Meta-LUN Name: LUN30
Meta-LUN WWN: 60:06:01:6D:35:0A:00:00:FC:35:5B:8C:66:60:D7:11
Meta-LUN Number: 30
Default Owner: SP A
Current Owner: SP B
```

```

Meta-LUN State:  ENABLED
Can MetaLUN Be Expanded:  Yes
Expansion Rate:  MEDIUM
Percent Expanded:  0
Total Capacity (Blocks/Megabytes):  819200/400
Actual User Capacity (Blocks/Megabytes):  819200/400
Autoassign:  No
Is Redundant:  Yes
Drive Type:  Fibre Channel
Element Size Multiplier:  40
Bind Offset:  0
Read Histogram [0] :    22
Read Histogram [1] :    14
Read Histogram [2] :    19
Read Histogram [3] :    37
Read Histogram [4] :     6
Read Histogram [5]:     0
Read Histogram [6]:     0
Read Histogram [7]:     0
Read Histogram [8]:     0
Read Histogram [9]:     0
Read Histogram Overflows:          45
Write Histogram [0] :    22
Write Histogram [1] :    22
Write Histogram [2] :    22
Write Histogram [3] :    22
Write Histogram [4] :    22
Write Histogram [5] :    22
Write Histogram [6] :    22
Write Histogram [7] :    22
Write Histogram [8] :    22
Write Histogram [9] :    22
Write Histogram Overflows:          45
Read Requests:  0
Write Requests: 0
Blocks Read:    0
Blocks Written: 0
Components:
Number of LUNs: 2
LUNs:
LUN Number:  513
LUN Name:    LUN 30
LUN WWN:    60:06:01:72:35:0A:00:00:B0:44:01:4B:61:60:D7:11
RAID Type:  RAID5
LUN Number:  512
LUN Name:    LUN 31
LUN WWN:    60:06:01:72:35:0A:00:00:15:FB:3B:4C:61:60:D7:11
RAID Type:  RAID5
Number of LUNs:  1 (Lists LUN information for second component)

```

The CLI returns the status or the error information if the command fails.

This command displays the active/active counters for this metaLUN along with other metaLUN output.

```
naviseccli -h 10.14.85.50 -user a -password a -scope 0 metalun -list
-metalunnumber 4 -aac
```

```

Host Blocks Read SPA:  6 (optimal)
Host Blocks Read SPB:  0
Host Blocks Written SPA:  0 (optimal)
Host Blocks Written SPB:  0
Host Read Requests SPA:  6 (optimal)
Host Read Requests SPB:  0
Host Write Requests SPA:  0 (optimal)

```

```

Host Write Requests SPB: 0
Busy Ticks SPA: 0 (optimal)
Busy Ticks SPB: 0
Idle Ticks SPA: 33224 (optimal)
Idle Ticks SPB: 0
Explicit Trespasses SPA: 0
Explicit Trespasses SPB: 0
Implicit Trespasses SPA: 0
Implicit Trespasses SPB: 0

```

metalun -modify

Modifies certain properties of an existing metaLUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli metalun` command with the `-modify` function lets you edit certain properties of a specified metaLUN. You must specify either the number or WWN for the metaLUN, and you must specify at least one other switch.

SYNTAX

```

metalun -modify -metalun number|WWN [-autoassign 0|1]
[-cap capacity] [-defaultowner A|B]
[-expansionrate low|medium|high|ASAP|value]
[-newname newName] [-o] [-sq tb|gb|mb|bc] [-tomaxcap]

```

OPTIONS

`-metalun number|WWN`

Specifies the number or WWN of the metaLUN whose properties you want to change.

`-autoassign 1|0` (Not supported on AX4-5 series systems)

Enables or disables auto assign for a metaLUN. Specify 1 to enable auto assign and 0 to disable it. If you do not specify an auto assign value, it defaults to the value of the primary LUN. For more information about auto assign, see the `-autoassign 0|1` option in [metalun -expand on page 184](#).

`-cap capacity`

Specifies the capacity of usable space in the metaLUN. The capacity must be smaller than the maximum capacity of the metaLUN, and greater than or equal to the current capacity (user capacity of the LUN or metaLUN before adding LUNs).

Note

If you want to modify the user capacity of a LUN or metaLUN that is participating in a SnapView (snapshot or clone), MirrorView or SAN Copy session, see [Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN on page 197](#).

`-defaultowner A|B`

Specifies the new default SP owner for the metaLUN — A or B.

`-expansionrate low|medium|high|ASAP|value`

Specifies the relative importance of making additional LUN capacity available to the host. It determines the amount of resources the SP devotes to re-striping the existing data on the original LUNs across the original LUNs and the new LUNs instead of to normal I/O activity. This rate is available only when you first create a metaLUN and is important only for a stripe expansion. Valid rebuild priorities are ASAP (default) or 3, High or 2, Medium or 1, and Low, or 0.

An expansion rate of ASAP or High re-stripes the data across the metaLUN faster than one with Medium or Low priority, but may degrade storage-system performance.

-newname *newName*

Specifies the new name for the metaLUN. A valid name must not exceed 64 characters and must include at least one non-whitespace character. This name must be unique on the system.

-o

Executes the command without prompting for confirmation.

-sq *tb|gb|mb|bc*

Specifies the size qualifier for the capacity of the metaLUN component.

where:

tb equals terabytes, gb equals gigabytes, mb equals megabytes and bc equals block count. The default qualifier is block count.

-tomaxcap

Lets you expand the metaLUN to its maximum capacity.

Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN

You must perform the operations detailed in the following sections to expand the capacity of a LUN or metaLUN that is currently participating in a SnapView, MirrorView or SAN Copy session (see [metalun -destroy on page 184](#)), and to allow the host to access the added capacity.

Note

For a stripe expansion, any existing data is re-striped across the original LUN or metaLUN and the added LUNs during the SnapView, MirrorView or SAN Copy session. This re-striping is invisible to the host.

For LUNs in a SnapView session

To access the added capacity of an expanded snapshot source LUN:

Procedure

1. Stop all SnapView sessions and destroy all snapshots associated with the expanded source LUN (see the SnapView Command Line Interface (CLI) Reference and the online help).
2. Use the [metalun -modify on page 196](#) command to change the capacity of the expanded LUNs.
3. Make sure the host operating system can utilize the added capacity.

Results

You can now start new sessions and create new snapshots with the expanded LUNs (see the *SnapView Command Line Interface (CLI) Reference* and the online help).

For LUNs in a clone group

Note

You must expand the clone and the clone source LUN to the same size.

To access the added capacity of an expanded clone, or clone source LUN, do the following:

Procedure

1. For the clone, remove the clone from the clone group; for a clone source LUN, destroy the clone group.
2. Use the [metalun -modify on page 196](#) command to change the capacity of the expanded LUNs.
3. Make sure the host operating system can utilize the added capacity.
4. Reconstruct the clone group with the expanded LUNs and then perform a full synchronization of the LUNs in the clone group.

For LUNs in mirrors

Note

You must expand both the primary and secondary images of a mirror to the same size.

To access the added capacity of an expanded primary and secondary MirrorView LUN, do the following:

Procedure

1. For the expanded primary image, destroy the mirror; for the expanded secondary image, remove it from the mirror (see the *MirrorView CLI Reference* or the online help).
2. Use the [metalun -modify on page 196](#) command to change the capacity of the expanded LUNs.
3. Make sure the host operating system can utilize the added capacity.
4. Reconstruct the mirror with the expanded LUNs and then perform a full synchronization of the mirror.

For LUNs in SAN Copy sessions

Note

If you expand the SAN Copy source logical unit, make sure that the destination LUNs are of equal or greater capacity.

To access the added capacity of an expanded SAN Copy source LUN or destination LUN, do the following:

Procedure

1. For any sessions that the expanded LUNs are participating in, wait until the SAN copy sessions complete, or stop the sessions, and then remove these sessions. (See the online help).
2. Use the [metalun -modify on page 196](#) command to change the capacity of the expanded LUNs.
3. Make sure the host operating system can utilize the added capacity.
4. Reconstruct the SAN Copy sessions that you removed and then, when you are ready, start the sessions. (See the online help.)

Reserved LUN Pool commands

This section explains how to configure the reserved LUN pool, which is required for running replication software such as SnapView, incremental SAN Copy, or MirrorView/

Asynchronous. It also describes the CLI commands that you can use if SnapView is installed on your system.

reserved -lunpool -addlun

Adds a LUN to the reserved LUN pool.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli reserved` command with the `-lunpool` and `-addlun` functions adds one or more LUNs to the reserved LUN pool.

Note

The reserved `-addlun` command restricts the Reserved LUN Pool LUNs to 128 element size. If you select LUN(s) with element size not equal to 128, the system will display the following error message: `Invalid element size for chosen LUN(s). Please choose LUN(s) with element size of 128.`

Each system has its own reserved LUN pool, and before starting a replication session, the reserved LUN pool must contain at least one LUN for each source LUN that will be participating in a session.

Note

This command does not support the thin LUNs. Replication sessions include SnapView sessions and any reserved sessions for use in another application, such as incremental SAN Copy and MirrorView/A.

The software allocates reserved LUNs on a per source LUN basis. Each system manages its own reserved LUN pool space and allocates reserved LUNs on a per-source LUN basis, not a per-session basis.

Note

You must bind the LUN before you can add it to the reserved LUN pool. While a LUN is part of the reserved LUN pool, you cannot use it for any other purpose.

SYNTAX

```
reserved -lunpool -addlun lun_numbers
```

OPTIONS

lun_numbers

Specifies the logical unit number(s) to add to the reserved LUN pool.

EXAMPLE

For `ss1_spa`, this command adds the LUN with ID 11 to the system's reserved LUN pool.

```
naviseccli -h ss1_spa reserved -lunpool -addlun 11
```

OUTPUT

If the version of software running on the SP does not support this command, the following error message is printed to `stderr`: `Command is not supported`

reserved -lunpool -list

Displays reserved LUN pool information.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli reserved` command with the `-lunpool` and `-list` functions and no other switches lists all information about the reserved LUN pool. You can obtain more specific information with function switches.

Note

This command does not support the thin LUNs.

SYNTAX

```
reserved -lunpool -list [-allocatedluns] [-chunksize]
[-freeluncount] [-freeluns] [-freesize] [-luncount]
[-percentused] [-size] [-used]
```

OPTIONS

`-allocatedluns`

Displays the LUN IDs of all allocated LUNs in the reserved LUN pool.

`-chunksize`

Displays the number of disk blocks in the reserved LUN pool, in KB. The chunk size applies to the entire system.

`-freeluncount`

Displays the total number of unallocated LUNs in the reserved LUN pool.

`-freeluns`

Displays the LUN IDs of all unallocated LUNs in the reserved LUN pool.

`-freesize`

Displays the total size of unallocated LUNs in GBs.

`-luncount`

Displays the total number of LUNs in the reserved LUN pool. You can add LUNs with the `reserved -lunpool -addlun` functions.

`-percentused`

Displays the percentage of the reserved LUN pool that is used.

`-size`

Displays the size of the reserved LUN pool in GBs.

`-used`

Displays the space used in the reserved LUN pool, in GBs.

EXAMPLE

This command displays the system's reserved LUN pool information. A sample output follows.

```
naviseccli -h ss1_spa reserved -lunpool -list
```

OUTPUT

The following output samples show devices controlled by one agent. Actual output varies depending on the command switches you use with the `-lunpool -list` command.

Note

If your system is running VNX CLI for Block version 6.24 or later, reserved LUN pool information will be displayed for the entire system, regardless of the SP you direct the command to. If your system is running VNX CLI for Block version 6.22 or earlier, reserved LUN pool information will be displayed for the SP you direct the command to.

`naviseccli output:`

```
Name of the SP: GLOBAL
Total Number of LUNs in Pool: 6
Number of Unallocated LUNs in Pool: 4
Unallocated LUNs: 14, 24, 23, 22
Total size in GB: 4.488281
Unallocated size in GB: 1.292969
Used LUN Pool in GB: 0.017639
% Used of LUN Pool: 0.393005
Chunk size in disk blocks: 128
Allocated LUN Pool:
Target LUN: 2
Allocated LUNs: 249
Lun Pool LUN % Used: 91.852825
Allocated LUNs: 250
Lun Pool LUN % Used: 0
```

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

reserved -lunpool -rmlun

Removes one or more LUNs from the reserved LUN pool.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli reserved` command with the `-lunpool` and `-rmlun` functions removes one or more LUNs from the reserved LUN pool. If you omit the `override` option, the CLI prompts for confirmation before removing the LUN from reserved LUN pool.

Note

This command does not support the thin LUNs.

SYNTAX

```
reserved -lunpool -rmlun luns [-o]
```

OPTIONS

`-o`

Executes the command without prompting for confirmation.

`-rmlun luns`

Specifies the LUN ID of each LUN, with the IDs separated by blanks, you want to remove from the reserved LUN pool.

EXAMPLE

For `ss1_spa`, this command starts removing LUN 42 from the system's reserved LUN pool. The software asks for confirmation:

```
naviseccli -h ssl_spa reserved-lunpool -rmlun 42
```

```
Do you really want to remove the luns from LUN pool.
(y/n) [N]? y
```

OUTPUT

If the version of software running on the SP does not support this command, a Command is not supported error message is printed to stderr.

iSCSI and FCoE commands

This section explains each of the navisecli Internet SCSI (iSCSI) and Fibre Channel over Ethernet (FCoE) commands, that is, the CLI commands that are common to all iSCSI systems and FCoE ports.

iSCSI commands overview

The navisecli iSCSI commands in this section include the connection and iSNS commands. [navisecli iSCSI commands on page 202](#) identifies the iSCSI commands that can be used.

Table 11 navisecli iSCSI commands

Command	Description
CHAP User Commands	
connection -adduser	Creates a CHAP user account.
connection -deleteuser	Deletes a CHAP user account.
connection -getuser	Displays CHAP user accounts.
Port administration commands	
connection -getport	Retrieves a specific iSCSI or FCOE port or the list of all the iSCSI or FCoE ports.
connection -setport	Modifies the properties of a connection port.
connection -delport	Deletes a virtual port.
connection -route	Sets up the static routes.
connection -pingnode	Verifies that packets reach a destination node without errors.
connection -traceroute	Displays the route that a packet takes to a destination node.
Connection set database commands	
connection -setsharedauth	Sets a username and secret for the system's shared credentials record.
connection -delsharedauth	Removes the shared authentication information from the system.
connection -getsharedauth	Displays the username and the switch enabled for the shared authentication record.
connection -addset	Adds a connection set.

Table 11 naviseccli iSCSI commands (continued)

Command	Description
<code>connection -delset</code>	Deletes a connection set and associated connection paths.
<code>connection -modifyset</code>	Modifies parameters of the connection set.
<code>connection -getset</code>	Lists a specific connection set or all the connection sets.
<code>connection -addpath</code>	Adds a connection path to a specific connection set.
<code>connection -delpath</code>	Deletes a connection path from a connection set.
<code>connection -modifypath</code>	Modifies a connection path from a connection set.
<code>connection -verifypath</code>	Tests a connection path from a connection set.
iSNS commands	
<code>isns -addserver</code>	Adds an iSNS server to a system's iSNS service configuration.
<code>isns -deleteserver</code>	Deletes an iSNS server from a system's iSNS service configuration.
<code>isns -listserver</code>	Displays all iSNS servers that are configured in the system's iSNS service.
<code>isns -setprimary</code>	Sets a configured iSNS server as the primary server for a system's iSNS service configuration.

connection -adduser

Adds a CHAP authentication user.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-adduser` function adds either an initiator CHAP authentication user or a target (system) CHAP authentication user to the system CHAP user database.

Note

If you try to create a CHAP user for an initiator and the user is already defined, an error message is returned that the user is already defined for the initiator.

If you try to add a user for a target and the user is already defined, an error is returned that only one user can be defined for target access.

SYNTAX

```
connection -adduser -definedFor initiator|target |
-initiatorName initiatorName | ANY [-userName userName]
-secret secret [-hexsecret] [-o]
```

OPTIONS

`-definedFor initiator | target`

initiator adds an initiator CHAP user account to the system user database. If you specify initiator, you must specify an initiator using the `-initiatorName` parameter. Target specifies that a target (system) CHAP user account will be created. If you specify target, do not use the `-initiatorName` parameter.

`-initiatorName` *initiatorName*

If initiator user account data is being added, `initiatorName` is a unique identifier for the initiator and uses the iSCSI Qualified Name (IQN) of the iSCSI initiator. This parameter is required if you are defining an initiator user.

`-userName` *username*

If initiator user account data is being added, `username` specifies the CHAP username that the initiator uses when connecting to an iSCSI target. For CHAP, a username and secret must be defined for each initiator. If the `-userName` parameter is omitted, it defaults to the `-initiatorName` value. If a target is being configured for mutual CHAP, the target name becomes the username.

`-secret` *secret*

If initiator user account data is being added, `secret` specifies the CHAP secret (password) that the system will use to authenticate the initiator. You can specify the secret in ASCII or hexadecimal characters (see the `-hexsecret` switch that follows). By default, `-secret` is an ASCII string. ASCII secrets must consist of at least 12 and no more than 16 characters. Hexadecimal secrets must consist of at least 12 and no more than 16 pairs of data (24 to 32 characters).

`-hexsecret`

If initiator user account data is being added, including the `-hexsecret` parameter specifies that `-secret` is a hexadecimal value.

`-o`

Executes the command without prompting for confirmation. If you include the `-o` switch, you will not be prompted to confirm that you want to add the user account data. If you do not use the `-o` switch, you will be prompted to confirm that you want to add the user.

EXAMPLE

The following example adds a user account to the system database.

```
naviseccli connection -adduser -definedfor initiator -initiatorName
5.com.microsoft:cpc7745 -username Guest -secret 1234567890123456
```

```
Adding an iSCSI user account with the following
attributes:
Initiator Name: iqn.1991-05.com.microsoft:cpc7745
User name: Guest
Secret: 1234567890123456
Defined For: initiator
Do you really want to perform this action (y/n)? y
```

connection -deleteuser

Deletes a CHAP authentication user.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-deleteuser` function deletes a specified initiator or target CHAP authentication user from the system CHAP user database.

SYNTAX

```
connection -deleteuser -definedFor initiator | target
-initiatorName initiatorName -userName username [-o]
```

OPTIONS

-definedFor initiator | target

initiator specifies that an initiator user account will be deleted from the system user database. If you specify initiator, you must specify an initiator using the **-initiatorName** parameter. target specifies that the target (system) user data will be deleted.

-initiatorName *initiatorName*

If initiator user account data is being deleted, *initiatorName* is the name that was assigned for the initiator when the user account was created. The default initiator name is the IQN of the initiator.

-userName *username*

Specifies the iSCSI CHAP username for the initiator or target user being deleted.

-o

Executes the command without prompting for confirmation. If you include the **-o** switch, you will not be prompted to confirm the user account deletion. If you do not use the **-o** switch, you will be prompted to confirm that you want to delete the user account data.

EXAMPLE

The following example deletes a user account from the system:

```
naviseccli connection -deleteuser -definedfor initiator -initiatorName
iqn.1991-05.com.microsoft:cpc7745 -username Guest
```

```
Deleting an iSCSI user account with the following
attributes:
Initiator Name: iqn.1991-05.com.microsoft:cpc7745
User name: Guest
Defined For: initiator
Do you really want to perform this action (y/n)? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -getuser

Display CHAP authentication user data.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-getuser` function displays CHAP authentication user account data that is configured on the system. If no arguments are specified, all CHAP users defined on the system are displayed. Use `-definedFor`, `-initiatorName`, and `-userName` to display information for a specified user. If you specify a particular user and that user is not found, the error message, `The specified user account does not exist`, is returned.

SYNTAX

```
connection -getuser [-definedFor initiator | target]
[-userName userName] [-initiatorName initiatorName]
```

OPTIONS

`-definedFor initiator | target`

`initiator` specifies that information for a CHAP initiator user is displayed. If you specify `initiator`, you must include the `-userName` and `-initiatorName`.

`target` specifies that information for a CHAP target user (the system user account) is displayed. If you specify `target`, you must include the `-userName`.

`-initiatorName initiatorName`

The initiator is the host or HBA that initiates the connection with the target system. The initiator is identified by the IQN of the iSCSI initiator.

`-userName username`

Specifies the iSCSI CHAP username for the initiator or target user being specified by `initiatorName`.

EXAMPLE

The following example gets information on the current user:

```
naviseccli connection -getuser
```

OUTPUT

```
Initiator Name:  iqn.1991-05.com.microsoft:cpc7745
User Name:  Guest
Defined For:  initiator
```

connection -getport

Retrieve a specific iSCSI or FCoE port or the list of all the iSCSI or FCoE ports.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-getport` function displays information about either a specified iSCSI port or a list of all iSCSI ports. If no arguments are specified, all of the iSCSI ports defined on the targeted system are displayed.

If a change in settings is requested, a confirmation is displayed.

The `-sp` and the `-portid` parameters select which iSCSI ports are displayed. If both options are specified, a particular iSCSI port is specified. If either or both of the options are omitted, all iSCSI ports that meet the search criteria are displayed.

SYNTAX

```
connection -getport [-sp a|b] [-portid portid [-vportid
vportid]] [-vlanid] [-address IPAddress] [-subnetmask]
[-gateway]
[-initiatorAuthentication] [-mac] [-mtu] [-flowctl]
[-hostWindow] [-replicationWindow] [-all]
```

OPTIONS

`-sp a|b`

Specifies SP A or B. Displays the information for iSCSI or FCoE ports configured on the specified SP.

`-portid portid`

Specifies an iSCSI or FCoE port ID. Displays information only for the specified iSCSI port. The following parameters, `-address`, `-subnetmask`, `-gateway`, and `-initiatorAuthentication`, define additional information that is displayed for

iSCSI ports. If none of these options is specified, all fields are displayed. The SP and port ID fields are always displayed.

`-address`

Displays the IP address of the specified iSCSI ports. The address is a 32-bit numeric address written as four 8-bit numbers (0-255) separated by periods (called the quad-dotted notation). For example, 165.152.36.35.

`-subnetmask`

Displays the network mask for the specified iSCSI ports. The subnet mask uses the quad-dotted notation. For example, 255.255.255.0.

`-gateway`

Displays the network gateway address that the specified iSCSI ports use. The gateway address uses the quad-dotted notation. For example, 165.152.35.22.

`-initiatorAuthentication`

Displays the state of the initiator authentication flag for the specified iSCSI ports. True indicates that initiator authentication is turned on for the port and that all initiators logging into that port will be required to authenticate against the target system. False indicates that initiator authentication has been disabled for the port.

`-mac`

Displays the MAC address for the front-end ports of an iSCSI system and also FCoE ports. Running the `port ... -mac` command returns a *MAC Address* field for each port. When you run the `port ... -mac` command, if any ports are Fibre Channel ports, the *MAC Address* field for those ports will be displayed as *Not Applicable*.

`-mtu`

Displays the port's MTU settings and lists the valid MTU settings for the portal.

`-vlanid`

Displays the virtual LAN (VLAN) ID of the specified virtual port.

`-vportid vportid`

Displays information for the specified virtual port ID.

`-flowctl`

Determines the flow control settings. Valid values are Auto, Transmit flow control (Tx), Receive flow control (Rx).

`-hostWindow`

Sets the receive window size for host-to-system traffic. The available window sizes are 64K, 128K, 256K, 512K, 1MB. The default value is 256K.

`replicationWindow`

Sets the receive window size for system-to-system traffic. The available window sizes are 64K, 128K, 256K, 512K, 1MB. The default value is 256K.

`-all`

Displays all the physical and virtual port information for each iSCSI port.

EXAMPLE #1

`naviseccli connection -getport`

```
SP: A
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a4
iSCSI Alias: 0877.a4
IP Address: 172.20.1.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: A
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a5
iSCSI Alias: 0877.a5

SP: A
```

```

Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a6
iSCSI Alias: 0877.a6
IP Address: 172.20.2.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: A
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a7
iSCSI Alias: 0877.a7

SP: B
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b4
iSCSI Alias: 0877.b4
IP Address: 172.20.1.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: B
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b5
iSCSI Alias: 0877.b5

SP: B
Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b6
iSCSI Alias: 0877.b6
IP Address: 172.20.2.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: B
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b7
iSCSI Alias: 0877.b7

SP: B
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:69:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

SP: A
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:61:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

```



```

SP: B
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:68:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

```

EXAMPLE #2

```
naviseccli connection -getport -sp a -portid 8
```

```

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:68:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

```

EXAMPLE #3

```
naviseccli connection -getport -sp a -portid 8 -vlanid
```

```

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
Virtual Port ID: 0
VLAN ID: 101

```

EXAMPLE #4

```
naviseccli connection -getport -sp a -portid 8 -mtu
```

```

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
Current MTU: 2240
Available MTU Sizes: "2240"?

```

EXAMPLE #5

```
naviseccli connection -getport -portid 8 -speed
```

```

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb

```

EXAMPLE #6

```
naviseccli connection -getport -sp a -portid 8 -vportid 0 -all
```

```
SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: 101
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
```

EXAMPLE #7

```
naviseccli connection -getport -mac
```

```
SP: A
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a4
iSCSI Alias: 0877.a4
Port Speed: 1000 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
-                : 100 Mb
-                : 1000 Mb
-                : Auto
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.1.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: A
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a5
iSCSI Alias: 0877.a5
Port Speed: 0 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
-                : 100 Mb
-                : 1000 Mb
-                : Auto
Current MTU: 0
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: A
Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a6
iSCSI Alias: 0877.a6
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
```

```

Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.2.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: A
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a7
iSCSI Alias: 0877.a7
Port Speed: 0 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b4
iSCSI Alias: 0877.b4
Port Speed: 1000 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
- : 100 Mb
- : 1000 Mb
- : Auto
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.1.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: B
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b5
iSCSI Alias: 0877.b5
Port Speed: 0 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
- : 100 Mb
- : 1000 Mb
- : Auto
Current MTU: 0
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b6
iSCSI Alias: 0877.b6
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb

```

```

Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.2.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: B
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b7
iSCSI Alias: 0877.b7
Port Speed: 0 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:69:3B:24:13:0D
iSCSI Alias: N/A
Enode MAC Address: 00-60-16-3B-47-B3
Port Speed: 10 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: 666
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 0E-FC-00-68-00-03

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
Enode MAC Address: 00-60-16-3B-55-42
Port Speed: 0 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 00-00-00-00-00-00

SP: A
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:61:3B:24:13:0D
iSCSI Alias: N/A
Enode MAC Address: 00-60-16-3B-55-43
Port Speed: 10 Gb

```

```

Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: 666
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 0E-FC-00-68-00-04

SP: B
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:68:3B:24:13:0D
iSCSI Alias: N/A
Enode MAC Address: 00-60-16-3B-47-B2
Port Speed: 0 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 00-00-00-00-00-00

```

EXAMPLE #8

```
naviseccli connection -getport -sp a -portid 8 -vportid 0 -mac
```

```

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Enode MAC Address: 00:60:16:32:12:5B

Virtual Port ID: 0
VLAN ID: 101
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 0E:FC:00:2B:0B:04

```

EXAMPLE #9

```
naviseccli -h peregrine-spa connection -getport -sp a -mac -all
```

```

SP: A
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:61:3B:24:13:0D
iSCSI Alias: N/A
Enode MAC Address: 00-60-16-3B-55-43
Port Speed: 10 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

```

```

Virtual Port ID: 0
VLAN ID: 666
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A
VN_Port MAC Address: 0E-FC-00-68-00-04

```

EXAMPLE #10

```

naviseccli -h peregrine-spa connection -getport -flowctl -hostWindow
-replicationWindow

```

```

Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

```

OUTPUT

The following sample output shows several virtual ports assigned with the VLAN tags:

```

naviseccli -h peregrine_spa connection -getport -sp a -portid 0 -all

```

```

SP: A
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a4
iSCSI Alias: 0877.a4
Port Speed: 1000 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
-                : 100 Mb
-                : 1000 Mb
-                : Auto
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.1.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: A
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a5
iSCSI Alias: 0877.a5
Port Speed: 0 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
-                : 100 Mb
-                : 1000 Mb
-                : Auto
Current MTU: 0
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

```

```

SP: A
Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a6
iSCSI Alias: 0877.a6
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.2.140
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: A
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.a7
iSCSI Alias: 0877.a7
Port Speed: 0 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 4
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b4
iSCSI Alias: 0877.b4
Port Speed: 1000 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
- : 100 Mb
- : 1000 Mb
- : Auto
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.1.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.1.1
Initiator Authentication: Not Available

SP: B
Port ID: 5
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b5
iSCSI Alias: 0877.b5
Port Speed: 0 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb

```

```

-           : 100 Mb
-           : 1000 Mb
-           : Auto
Current MTU: 0
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 6
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b6
iSCSI Alias: 0877.b6
Port Speed: 10000 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: Disabled
IP Address: 172.20.2.141
Subnet Mask: 255.255.255.0
Gateway Address: 172.20.2.1
Initiator Authentication: Not Available

SP: B
Port ID: 7
Port WWN: iqn.1992-04.com.emc:cx.apm00093300877.b7
iSCSI Alias: 0877.b7
Port Speed: 0 Mb
Auto-Negotiate: No
Available Speeds: 10000 Mb
Current MTU: 1500
Available MTU Sizes:
"1260","1448","1500","1548","2000","2450","3000","4000","4080","4470",
"5000","6000","7000","8000","9000"

SP: B
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:69:3B:24:13:0D
iSCSI Alias: N/A
Port Speed: 10 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: 666
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

SP: A
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:60:3B:24:13:0D
iSCSI Alias: N/A

```



```

Port Speed: 0 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

SP: A
Port ID: 9
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:61:3B:24:13:0D
iSCSI Alias: N/A
Port Speed: 10 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"
Flow Control: Auto
Host Window: 256K
Replication Window: 256K
Available Window Sizes: 64K,128K,256K,512K,1MB

Virtual Port ID: 0
VLAN ID: 666
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

SP: B
Port ID: 8
Port WWN: 50:06:01:60:BB:20:13:0D:50:06:01:68:3B:24:13:0D
iSCSI Alias: N/A
Port Speed: 0 Gb
Auto-Negotiate: No
Available Speeds: 10 Gb
Current MTU: 2240
Available MTU Sizes: "2240"

Virtual Port ID: 0
VLAN ID: N/A
IP Address: N/A
Subnet Mask: N/A
Gateway Address: N/A
Initiator Authentication: N/A

```

connection -setport

Modify the properties of a port.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-setport` function modifies the properties of a specified SP port.

SYNTAX

```
connection -setport [-iscsi (Default)]
-sp a|b -portid portid [-vportid vportid (default=0)]
[-vlanid vlanid|-vlandisable]
[-address address] [-subnetmask subnetmask] [-gateway gateway]
[-initiatorAuthentication 0|1] [-alias iscsialias]
[-speed speed] [-mtu mtu] [-o]
```

```
connection -setport [-fcoe]
-sp a|b -portid portid [-vportid vportid (default=0)]
[-vlanid vlanid|-vlanautodiscovery] [-o]
```

OPTIONS

-sp a|b

Specifies that the iSCSI or FCOE port is on either SP A or SP B.

-portid *portid*

Specifies the port on the SP. For example, an SP with one port has port 1 and an SP with two ports has port 0 and port 1.

-vportid *vportid*

Specifies the virtual port ID. If you do not specify the **-vportid**, the system defaults to virtual port 0.

-vlanautodiscovery

Sets the port (FCoE only) specified by **-portid** to be in VLAN auto discovery mode.

-vlanid *vlanid*

Specifies the VLAN ID.

-vlandisable

Disables VLAN tagging on the virtual port.

-address *address*

Specifies the IP address of the iSCSI port. The address uses the dotted-quad format. For example: 165.152.36.35.

-subnetmask *subnetmask*

Specifies the subnetwork mask for the iSCSI port. The subnet mask uses the dotted-quad format. For example: 255.255.255.0.

-gateway *gateway*

Specifies the network gateway address that the iSCSI port uses. The gateway address uses the dotted-quad format. For example: 165.152.35.22.

-initiatorAuthentication 0|1

If initiator authentication is turned on, all initiators logging in to the port must be authenticated by the system.

0 = Not required - default

1 = Required

-alias *iscsialias*

Specifies the iSCSI alias.

-speed *speed*

Specifies the port speed.

-mtu

Configures the MTU size for the specified iSCSI port. If not specified during the initial creation of a portal, the default is 1500. If **-mtu** is not specified during the update of a portal's setting, the MTU setting will not change. If a change in settings has been requested, a confirmation will be displayed.

-o

Executes the command without prompting for confirmation. If you include the `-o` switch, the port will be set (configured) without a confirmation prompt. Without the `-o` switch, the software displays confirmation messages that you must respond to.

EXAMPLE #1

The following example configures the network settings for the port with IP address 10.14.80.110:

```
naviseccli connection -setport -sp b -portid 0 -address -subnetmask
255.255.255.0 -gateway 10.14.80.1
```

It is recommended that you consult with your Network Manager to determine the correct settings before applying these changes. Changing the port properties may disrupt traffic on one or more ports of this I/O Module. Initiator configuration changes may be necessary to regain connections.
Do you really want to perform this action (y/n)?y

EXAMPLE #2

```
naviseccli connection -setport -hostWindow 32K -o
```

Request failed. Invalid window size: 32K.
Available Window Sizes: 64K, 128K, 256K, 512K, 1MB

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -delport

Deletes a virtual port.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-delport` function lets you delete a virtual port from a physical port. To execute this command, more than one virtual port should be associated with the management port.

If there is only one management port with only one virtual port, the command cannot be executed. Use of this command depends on the network provider capabilities.

SYNTAX

```
connection -delport -sp a|b -portid portid -vportid vportid
[-o]
```

OPTIONS

`-sp a|b`

Deletes the virtual port of the specified SP. The default is the connected SP.

`-portid portid`

Specifies the iSCSI port ID. If you specify an invalid port ID, the system displays the following error: Request failed. Specified iSCSI port not found.

`-vportid vportid`

Specifies the virtual port ID. If you specify an invalid virtual port ID, the system displays the following error message: Request failed. Specified iSCSI virtual port not found.

`-o`

Executes the command without prompting for confirmation.

OUTPUT

Provides output depending on the options used.

connection -route

Sets up the static routes to specific hosts or networks.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-route` function is used to set up the static routes to specific hosts or networks.

SYNTAX

```
connection -route [-sp a|b] [-portid portid|-vportid vportid]
[-add -destination destination [-subnetmask subnetmask|
-prefixlength length] [-metric metric] [-persistent]]
[-delete -destination destination [-subnetmask subnetmask|
-prefixlength length]]
[-list] [-o]
```

OPTIONS

`-sp a|b`

Specifies the destination SP. The default is the connected SP.

`-portid portid`

The `-portid` and the `-vportid` together specify the interface associated with the route. If you do not specify the interface, the system determines it.

`-vportid vportid`

The `-portid` and the `-vportid` together specify the interface associated with the route. If you do not specify the interface, the system determines it.

`-add`

Adds a network route.

`-destination destination`

Specifies the IPv4/IPv6 address or hostname of the destination.

`-subnetmask netmask`

Specifies the IPv4 subnet mask value for the route entry. The default is 255.255.255.255. It is valid only for an IPv4 destination address.

`-prefixlength length`

Specifies the prefix length, that is the part of the address to compare while determining the route. The default is 128. It is valid only for an IPv6 destination address.

`-metric metric`

Determines the best route using the routing algorithm. The route with the lowest metric is preferred over another route.

`-persistent`

Specifies whether the route should persist across reboots.

`-delete`

Deletes an existing route.

`-list`

Displays the routes for the SP. By default, it displays the routes for the connected SP. It also displays the routes for a specific port/virtual port.

-o

Executes the command without prompting for confirmation.

OUTPUT

Provides output depending on the options used.

connection -pingnode

Verifies that packets reach a destination node without errors.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-pingnode` function transmits packets to a target node and returns a notification that the packets have reached the node without errors. The `connection -pingnode` command pings a target node from the iSCSI port. Use the `connection -pingnode` and the `connection -traceroute` commands to verify and diagnose network connectivity.

SYNTAX

```
connection -pingnode [-sp a|b] [ [-portid portid -vportid
vportid] | -source sourceaddress] -address address
[-packetSize packetSize] [-count count] [-timeouts timeouts]
[-delays delays]
```

OPTIONS

`-sp a|b`

Specifies either storage processor A or storage processor B.

`-portid portid`

Specifies the port on the SP, either 0 or 1. For example, an SP with one port has port 0 and an SP with two ports has port 0 and port 1.

`-vportid vportid`

Specifies the virtual port ID.

`-address address`

Specifies the IP address of the target node that the iSCSI port will ping. The address is a 32-bit numeric address written as four numbers separated by periods, for example, 128.221.56.52. Network names are not supported.

`-packetSize packetSize`

Specifies the size of the packets in bytes. The default is 32 bytes, the minimum is 0 bytes, and the maximum is 2048 bytes.

`-count count`

Specifies the number of pings to send. The default is 4; the minimum is 1.

`-timeouts timeouts`

Specifies the response wait timeout period in seconds. The default is 3 seconds, the minimum is 1 second, and the maximum is 15 seconds.

`-delays delays`

Specifies a delay in seconds between pings. The default is 1 second, the minimum is 1 second, and the maximum is 10 seconds.

EXAMPLE

```
naviseccli -h peregrine_spa connection -pingnode -sp a -portid 0
-vportid 0 -address 128.222.132.100
```

OUTPUT

```
Reply from 128.222.132.100: bytes=32 time=1ms TTL=30
Reply from 128.222.132.100: bytes=32 time=1ms TTL=30
Reply from 128.222.132.100: bytes=32 time=1ms TTL=30
Reply from 128.222.132.100: bytes=32 time=1ms TTL=30
```

connection -traceroute

Display the route that a packet takes to a destination node.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-traceroute` function displays the route that packets take to a target node that you specify in the command. The `connection -traceroute` command traces and displays the route from the system iSCSI port to the target node. Use the `connection -traceroute` and `connection -pingnode` commands to verify and diagnose network connectivity.

SYNTAX

```
connection -traceroute[-sp a|b] -address IPAddress
[ [-portid portid -vportid vportid] | -source sourceaddress
[-mtu| -packetSize packetSize] ]
[-timeouts timeouts] [-delays delays]
```

OPTIONS

`-address IPAddress`

Specifies the IP address of the target node that the iSCSI port will attempt to contact. The address uses the dotted-quad format, for example, 128.221.56.52. Network names are not supported.

`-portid portid`

Specifies the port on the SP, either 0 or 1. For example, an SP with one port has port 0; an SP with two ports has port 0 and port 1.

`-vportid vportid`

Specifies the virtual port ID.

`-sp a|b`

Specifies either SP A or SP B.

`-delays delays`

Specifies a delay in seconds between datagrams. The default is 1 second, the minimum is 1 second, and the maximum is 10 seconds.

`-packetSize packetSize`

Specifies the size of the echo packets in bytes. The default is 32 bytes, the minimum is 0 bytes, and the maximum is 2048 bytes.

`-mtu`

Performs networking diagnostics, focusing on the maximum transmission unit between the two systems. If specified, `-source` is required.

`-source sourceaddress`

Specifies the source address and is required when `-mtu` is used.

`-timeouts timeouts`

Specifies the response wait timeout period in seconds. The default is 3 seconds, the minimum is 1 second, and the maximum is 15 seconds.

EXAMPLE

The following example displays the route taken by a packet:

```
naviseccli -user a -password a -scope 0 -h ssl_spa connection
-traceroute -sp a -address -mtu -source
```

OUTPUT

```
1: 172.20.1.1 1 ms 1500 mtu
2: 10.5.1.250 1 ms 1500 mtu
```

connection -setsharedauth

Sets username and secret for the system's shared credentials record.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-setsharedauth` function sets a username and secret for the iSCSI initiator system's shared credentials record of the connection set database. The command lets you enable or disable the use of the system's shared authentication credentials.

SYNTAX

```
connection -setsharedauth[-userNameuserName] [-secret secret
[-hexsecret]] [-enable|-disable] [-o]
```

OPTIONS

`-userName userName`

Specifies the CHAP username used by the initiator when challenged by the target system. For CHAP, a username and secret must be defined.

`-secret secret`

Specifies the CHAP secret to be used with the username when challenged by the target system. You can specify the secret in ASCII or hexadecimal characters (see the `-hexsecret` switch that follows). By default, `-secret` is an ASCII string. ASCII secrets must consist of at least 12 and no more than 16 characters. Hexadecimal secrets must consist of at least 12 and no more than 16 pairs of data (24 to 32 characters).

`-hexsecret`

If initiator user account data is being added, including the `-hexsecret` parameter, it specifies that `-secret` is a hexadecimal value.

`-enable|-disable`

Enables or disables the use of the system's shared credentials.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example sets a username and secret for the system's shared credentials record:

```
naviseccli connection -setsharedauth -username user1
-secret abcdefghi123 -enable -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -delsharedauth

Removes the system shared authentication credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-delsharedauth` function removes the username and secret from the iSCSI initiator system.

SYNTAX

```
connection -delsharedauth[-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation. If you include the `-o` switch, you will not be prompted to confirm the user account deletion. If you do not use the `-o` switch, you will be prompted to confirm that you want to delete the user account data.

EXAMPLE

The following example removes the shared authentication credentials from the iSCSI initiator system:

```
connection -delsharedauth -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -getsharedauth

Displays the username and status of the shared authentication for the iSCSI initiator system

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-getsharedauth` function displays the username and status of the shared authentication for the iSCSI initiator system.

SYNTAX

```
connection -getsharedauth
```

EXAMPLE

The following example displays the username and status of the shared authentication for the iSCSI initiator system:

```
connection -getsharedauth
```

OUTPUT

```
User Name: user1
Enabled: True
```


If no shared credentials are defined for the system, then the output is:

```
No system shared credentials are defined.
```

connection -addset

Adds a connection set.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-addset` function adds a connection set to the connection set database.

SYNTAX

```
connection -addset -name name
[-digest none|Header|Data|Both]
[-authpolicy none|shared|setspecific -userName userName
-secret secret [-hexsecret] ] [-o]
```

OPTIONS

`-name name`

Specifies a unique identifier for the connection set. It is an arbitrary name supplied by the administrator that indicates the specific connection set using the CLI or UI.

`-digest none | Header | Data | Both`

Indicates the type of digest, a data protection mechanism, which is turned on for connections established through this connection set.

Header indicates that the iSCSI packet headers have a checksum applied and verified.

Data indicates that the data portion of the iSCSI packet have a checksum applied and verified. Both indicates that both the headers and the data portion of the iSCSI packet have a checksum applied and verified.

`-authpolicy none | shared | setspecific`

Defines the type of authentication this connection set uses.

`none` specifies that no credentials are needed to connect to the target system.

`shared` specifies that the system shared authentication credentials are used to authenticate.

`setspecific` indicates that the username and secret stored as part of this connection set are used for authentication.

`-userName userName`

Specifies the CHAP username used by the initiator when challenged by the target system. For CHAP, a username and secret must be defined.

`-secret secret`

Specifies the CHAP secret to be used with the username when challenged by the target system. You can specify the secret in ASCII or hexadecimal characters (see the `-hexsecret` switch that follows). By default, `-secret` is an ASCII string. ASCII secrets must consist of at least 12 and no more than 16 characters. Hexadecimal secrets must consist of at least 12 and no more than 16 pairs of data (24 to 32 characters).

`-hexsecret`

If initiator user account data is being added, including the `-hexsecret` parameter specifies that `-secret` is a hexadecimal value.

-o

Executes the command without prompting for confirmation.

EXAMPLE

The following example adds a connection set to the connection set database:

```
connection -addset -name MirrorSet -authpolicy none -digest none -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -delset

Deletes a connection set and associated connection paths.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-delset` function deletes a connection set and its associated connection paths by the name of the connection set.

SYNTAX

```
connection -delset -name name [-o]
```

OPTIONS

`-name name`

Specifies a unique identifier for the connection set. It is an arbitrary name supplied by the administrator that indicates the specific connection set using the CLI or UI.

-o

Executes the command without prompting for confirmation. If you include the `-o` switch, you will not be prompted to confirm the user account deletion. If you do not use the `-o` switch, you will be prompted to confirm that you want to delete the user account data.

EXAMPLE

The following example deletes a connection set and associated connection paths using the name of the connection:

```
connection -delset -name MirrorSet -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -modifyset

Modifies parameters of a connection set.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-modifyset` function modifies parameters of a connection set.

Note

You cannot change the name of the connection with this command.

SYNTAX

```
connection -modifyset -name name [-digest None|Header|Data|Both] [-authpolicy none|shared|setspecific -userName userName -secret secret [-hexsecret] ] [-o]
```

OPTIONS

-name *name*

Specifies a unique identifier for the connection set.

-digest None | Header | Data | Both

Indicates the type of digest, which is turned on for connections established through this connection set.

Header indicates that the iSCSI packet headers have a checksum applied and verified.

Data indicates that the data portion of the iSCSI packet have a checksum applied and verified.

Both indicates that both the headers and the data portion of the iSCSI packet have a checksum applied and verified.

-authpolicy none | shared | setspecific

Defines the type of authentication this connection set uses.

none specifies that no credentials are needed to connect to the target system.

shared specifies that the system shared authentication credentials are used for authentication.

setspecific indicates that the username and secret stored as part of this connection set are used for authentication.

-userName *userName*

Specifies the CHAP username used by the initiator when challenged by the target system. For CHAP, a username and secret must be defined.

-secret *secret*

Specifies the CHAP secret to be used with the username when challenged by the target system. You can specify the secret in ASCII or hexadecimal characters (see the **-hexsecret** switch that follows). By default, **-secret** is an ASCII string. ASCII secrets must consist of at least 12 and no more than 16 characters. Hexadecimal secrets must consist of at least 12 and no more than 16 pairs of data (24 to 32 characters).

-hexsecret

If initiator user account data is being added, including the **-hexsecret** parameter specifies that **-secret** is a hexadecimal value.

-o

Executes the command without prompting for confirmation.

EXAMPLE

The following example modifies the parameters of a connection:

```
connection -modifyset -name MirrorSet -authpolicy -setspecific
-username user2 -secret Twelve2Sixteen -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -getset

Lists a specific connection set or all the connection sets

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-getset` function lists a specific connection set or all the connection sets and displays the contents of each set.

SYNTAX

```
connection -getset [-name name]
```

OPTIONS

`-name name`

Specifies a unique identifier for the connection set.

EXAMPLE

The following example lists a specific connection set:

```
connection -getset -name MirrorSet
```

OUTPUT

```
Connection Set Name: MirrorSet
User Name: username
Authentication Usage: SetSpecific
Header Digest: Disabled
Data Digest: Disabled
IP Address of Target: 192.168.0.1
Listening IP Port of Target: 3260
SP Ports: A-0 A-1 B-0 B-1
Description: any comments that were entered.
```

Sample output if the system (CX4 series) has virtual ports:

```
Connection Set Name: tryzub_generated
User Name: TRISHULA_MirrorViewGenerated
Authentication usage: Setspecific
Header Digest: Enabled
Data Digest: Enabled
IP Address of Target: 172.20.4.146
Listening IP port of Target: 3260
SP Ports: A-3v0 B-2v4
Description: Path to tryzub A-1, generated for MirrorView
Listening IP port of Target: 3260
SP Ports: B-3v1
Description: Path to tryzub B-1, generated for MirrorView
```

connection -addpath

Adds a connection path to a specific connection set.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-addpath` function adds a connection path to a specific connection set. A connection path consists of the IP address of the target system, front-end ports used, and a comment field.

SYNTAX

```
connection -addpath -name name -target address[:port] [-sp a|b
-portid portid [-vportid vportid]] [comment text] [-o]
```

OPTIONS

- name *name*
Specifies the name of the connection set to which you will add the connection path.
- target *address:port*
Specifies the IP address and listening port (default 3260) of the target system.
- sp *a|b*
Specifies the storage processor on the initiating system used for the session.
- portid *portid*
Specifies the front-end port number on the storage processor of the initiating system to be used for the session.
- vportid *vportid*
Specifies the virtual port number for the session. If you do not specify the virtual port number, then the system defaults to virtual port 0.
- comment *text*
Specifies a comment or description of the connection path.
- o
Executes the command without prompting for confirmation.

EXAMPLE

The following example adds a connection path to the specific connection set:

```
connection -addpath -name MirrorSet -target 123.231.1.1 -sp a
-portid 3 -comment mirrorview connection path -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -delpath

Deletes a connection path from a connection set.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-delpath` function deletes a connection path from a connection set.

Note

You must have an IP address and optional port number to complete this command.

SYNTAX

```
connection -delpath -name name -target address[:port] [-o]
```

OPTIONS

- name *name*
Specifies the name of the connection set, which includes the path you want to delete.
- target *address:port*
Specifies the IP address and listening port (default 3260) of the target system.
- o

Standard override option. Without this option, the user is prompted to confirm intent.

EXAMPLE

The following example deletes a connection path from a connection set:

```
connection -delpath -name MirrorSet -target 123.231.1.1 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -modifypath

Modifies the parameters of a connection path.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-modifypath` function modifies the parameters of a connection path.

Note

You must have an IP address and optional port number to complete this command.

SYNTAX

```
connection -modifypath -name name -target address[:port] -add
-sp a|b -portid portid [-vportid vportid] [-comment text]
[-o] |
-add -comment text [-o] | -del -sp a|b -portid
portid [-vportid vportid] [-o]
```

OPTIONS

`-name name`

Specifies the name of the connection set, which includes the path you want to modify.

`-target address:port`

Specifies the IP address and listening port (default 3260) of the target system.

`-sp a|b`

Specifies the storage processor on the initiating system used for the session.

`-portid portid`

Specifies the front-end port number on the storage processor of the initiating system to be used for the session.

`-vportid vportid`

Specifies the virtual port number to be used for the session. If the virtual port number is not specified, the system defaults to virtual port 0.

`-comment text`

Specifies a new comment added to the connection path.

`-o`

Standard override option. Without this option, the user is prompted to confirm intent.

EXAMPLE

The following example modifies a connection path from a connection set:

```
connection -modifypath -name MirrorSet -target 23.231.1.1 -add -sp
b -portid 3 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

connection -verifypath

Tests the parameters of a connection path.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli connection` command with the `-verifypath` function tests the parameter of a connection path.

Note

A connection set name, an IP address, and an optional port number are needed to complete this command.

SYNTAX

```
connection -verifypath -name name -target address[:port]
[-sp a|b -portid portid [-vportid vportid] ]
```

OPTIONS

`-name name`

Specifies the name of the connection set, which contains the path to verify.

`-target address:port`

Specifies the IP address and listening port (default 3260) of the target system.

`-sp a|b`

Specifies the storage processor on the initiating system used for the session.

`-portid portid`

Specifies the front-end port number on the storage processor of the initiating system to be used for the session.

`-vportid vportid`

Specifies the virtual port number to be used for the session. If the virtual port number is not specified, the system defaults to virtual port 0.

`-o`

Standard override option. Without this option, the user is prompted to confirm intent.

EXAMPLE

The following example tests a connection path from a connection set:

```
Connection -verifypath -name MirrorSet -target 23.231.1.1 -o
```

OUTPUT

```
Test/verify command for target {ipaddr[:portnum]} from
{fe port} completed successfully.
```

isns -addserver

Adds an iSNS server to a system's iSNS service configuration.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli isns` command with the `-addserver` function adds an iSNS server to a system's iSNS service configuration.

When you add a server, if it is the first server you add to a system's iSNS service configuration, it is designated as the primary server. If a server exists in the system's iSNS service configuration, any subsequent server you add is designated as a backup.

Note

The primary server is the server the iSNS service actively communicates with.

SYNTAX

```
isns -addserver IPAddr [-o]
```

OPTIONS

IPAddr

The IP address of the iSNS server you want to add.

Note

You must identify the iSNS server by the server's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example adds the specified iSNS server, to the system's iSNS service configuration:

```
naviseccli -address ssl_spa isns -addserver 111.222.33.44
```

```
Adding an iSNS server with the following IP address: 111.222.33.44.
Do you really want to perform this action (y/n)? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

isns -deleteserver

Deletes an iSNS server from a system's iSNS service configuration.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli isns` command with the `-deleteserver` function, deletes an iSNS server from a system's iSNS service configuration.

Note

If the IP address of the server you want to delete is not configured, you will receive an error message.

If you delete a primary server, the next server listed in the system's iSNS service configuration becomes the primary.

Note

The primary server is the server the iSNS service actively communicates with.

SYNTAX

```
isns -deleteserver IPAddr [-o]
```

OPTIONS*IPAddr*

The IP address of the iSNS server you want to delete.

Note

You must identify the iSNS server by the server's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

-o

Executes the command without prompting for confirmation.

EXAMPLE

The following example deletes the specified iSNS server from the system's iSNS configuration:

```
naviseccli -address ssl_spa isns -deleteserver 111.222.33.44
```

```
Deleting an iSNS server with the following IP address:
111.222.33.44.
Do you really want to perform this action (y/n)? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

isns -listserver

Displays all iSNS servers that are configured in the system's iSNS service.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli isns` command with the `-listserver` function, lists all servers that are configured in a system's iSNS service.

If multiple servers are configured in the system's iSNS service, the `-listserver` operation will list multiple IP addresses and server types. You can have only one primary server; therefore you will only have one server designated as the primary. All other servers are designated as backup servers.

Note

The primary server is the server the iSNS service actively communicates with.

SYNTAX

```
isns -listserver
```

EXAMPLE

The following example lists the primary and backup iSNS servers for the system:

```
naviseccli -address ssl_spa isns -listserver
```

OUTPUT

```
IP Address: 111.222.33.44
Server Type: Backup
IP Address: 111.222.33.45
Server Type: Primary
```

IP Address is the IP address of the iSNS server. Server Type specifies whether the server is set as the primary or backup server.

isns -setprimary

Sets a configured iSNS server as the primary server for a system's iSNS service configuration.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli isns` command with the `-setprimary` function, establishes a configured iSNS server as the primary server in a system's iSNS service configuration. The primary server is the server the iSNS service actively communicates with. You can have only one primary server.

Note

If the IP address of the server you specify to be the primary server is not configured, you will receive an error message.

Use this command when you want to change the primary server in a system's iSNS service configuration. You can also use this command to re-establish a primary server after a failover. When a failure occurs that takes the existing primary server out of service, the iSNS service automatically establishes a backup server as the new primary server. You can use the `isns -setprimary` command to restore the original primary server as the primary server.

Note

If you have more than one backup server, the first backup server listed becomes the new primary server (see `isns -listserver`).

SYNTAX

```
isns -setprimary IPAddr [-o]
```

OPTIONS

IPAddr

The IP address of the iSNS server you want to designate as the primary server.

Note

You must identify the iSNS server by the server's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

-o

Executes the command without prompting for confirmation.

EXAMPLE

The following example establishes the specified iSNS server as the primary server:

```
naviseccli -address ssl_spa isns -setprimary 111.222.33.44
```

```
The listed IP Address will become the Primary iSNS
server: 111.222.33.44.
Do you really want to perform this action (y/n)? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

LUN Migration commands

This section explains the VNX CLI for Block commands that manage LUN migration.

LUN migration lets you improve the performance of a LUN by migrating the data from the source LUN to a destination LUN that has more desirable performance characteristics.

migrate -start

Starts a LUN migration session.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli migrate` command with the `-start` switch starts a LUN migration session.

Note

If the maximum number of migrations already exists on the subsystem, the new session is still created, but its state is queued.

For raw device mapping volumes (RDMS), if you are running VMware ESX Server and you are using the migration feature to copy LUNs to larger LUNs only, after you complete the migration you must rescan the bus at the ESX Server level. You can also rescan at the virtual machine level (optional). You then must expand the file system at the virtual machine level. ESX Server and the virtual machine will now recognize the completed migration.

For configuration specifications, go to the E-Lab Interoperability Navigator on EMC Online Support.

SYNTAX

```
migrate -start -source number|WWN -dest number|WWN -rate low|
medium|high|ASAP|value [-destroySnapshots] [-o]
```

OPTIONS

`-source number|WWN`

Specifies the source LUN for the migration. You can specify the LUN ID or World Wide Name (WWN).

`-dest number|WWN`

Specifies the destination LUN of the migration. You can specify the LUN ID or World Wide Name.

Note

The destination LUN must be the same size or larger than the source LUN.

`-rate low|medium|high|ASAP|value`

Specifies the LUN migration rate. Valid rates are low or 0, medium or 1, high or 2, and ASAP or 3.

Note

Do not use the ASAP migration rate when the system is in production, as the normal host I/O may be interrupted. Use ASAP only when the system is offline (free of any host-level I/O).

`-destroySnapshots`

Deletes snapshots automatically after migration if source LUN has snapshots, including any snapshots that were taken of consistency groups while this LUN was a member.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command starts a LUN migration session where the source LUN's ID is 6, the destination LUN's ID is 7, and the LUN migration rate is low.

```
naviseccli migrate -start -source 6 -dest 7 -rate low
```

OUTPUT

None if the command succeeds; status or error information if it fails.

migrate -cancel

Cancels a LUN migration session.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli migrate` command with the `-cancel` switch cancels an in-process LUN migration. The destination LUN is deleted when canceling occurs. Canceling the migration does not cause any data loss. The original LUN remains unchanged.

Note

You can cancel only one migration at a time.

SYNTAX

```
migrate -cancel -source number|WWN [-o]
```

OPTIONS

`-source number|WWN`

Specifies the source LUN for migration as the LUN ID or World Wide Name (WWN).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command cancels a LUN migration session whose source LUN's ID is 6.

```
naviseccli migrate -cancel -source 6
```

OUTPUT

None if the command succeeds; status or error information if it fails.

migrate -modify

Modifies a LUN migration session.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli migrate` command with the `-modify` switch modifies a LUN migration session according to the parameters you specify.

SYNTAX

```
migrate -modify -source number|WWN [-rate low|medium|high|ASAP|value] [-o]
```

OPTIONS

`-source number|WWN`

Specifies the source LUN of the migration you want to modify. You can specify the LUN ID or World Wide Name (WWN).

`-rate low|medium|high|ASAP|value`

Specifies the LUN migration rate. Valid rates are low or 0, medium or 1, high or 2, and ASAP or 3.

Note

Do not use the ASAP migration rate when the system is in production, as the normal host I/O may be interrupted.

Use ASAP only when the system is offline (free of any host-level I/O).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command modifies the LUN migration rate to medium.

```
naviseccli migrate -modify -source 6 -rate medium
```

OUTPUT

None if the command succeeds; status or error information if it fails.

migrate -list

Lists LUN migration sessions and their properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli migrate` command with the `-list` switch lists the existing LUN migration sessions and their properties.

SYNTAX

```
migrate -list [-source number|WWN
[-dest]
[-rate]
[-state]
[-percentcomplete]
[-timeremaining]
```

OPTIONS

`-source number|WWN`

Specifies the source LUN of the migration. You can specify the LUN ID or World Wide Name (WWN). If you specify the `-source` switch with a LUN, the command lists properties for only that migration. If you do not specify a source LUN, the output lists the properties for all migrations.

`-dest`

Displays the destination LUN.

`-rate`

Displays the LUN migration rate as low, medium, high, or ASAP.

`-state`

Displays the migration state as queued, migrating, migrated, transitioning, or faulted.

`-percentcomplete`

Displays the percentage of the migration that has completed.

`-timeremaining`

Displays the estimated time remaining for the migration to complete. For the first poll cycle after the migration is started, the Time Remaining field is displayed as ? (question mark).

EXAMPLE

This command lists the migration sessions and their properties.

```
naviseccli migrate -list
```

OUTPUT

```
Source LU Name: LUN 6
Source LU ID: 6
Dest LU Name: LUN 7
Dest LU ID: 7
Migration Rate: LOW
Current State: MIGRATING
Percent Complete: 12
Time Remaining: 1 minute(s)
```

Domain and Security commands

This section provides a brief overview of the domain and security features, and includes a list of commands you can use to manage system domains and user access to systems.

About domains and user accounts

A domain is a group of one or more systems with Storage Management Server software whose SPs are connected to a network and which have been assigned to the domain by VNX CLI for Block or Unisphere. Each domain has a master node (master system) that maintains the master copy of the domain data — the systems and global user accounts that make up the domain.

Setting up a domain allows a group of systems to be monitored and managed using a single login. Even if you plan to use a system by itself (manage it separately), we suggest that you create a domain for that system.

Note

When you setup a user account, you assign a scope of either global, local, or LDAP. Global user accounts apply to all systems within a domain. Local user accounts apply to a specific system. LDAP uses an external server to authenticate accounts for an entire domain.

VNX CLI for Block versions 6.24 and later support lightweight directory access protocol (LDAP). This requires an active and operational LDAP environment to which you can connect Unisphere security for the purpose of sharing user accounts and authenticating usernames and passwords.

Administrators can manage accounts in a centralized directory accessible to multiple applications. Users can use their LDAP credentials to log in and perform Unisphere and CLI operations.

Note

In order to issue domain CLI commands, you must have global administrator privileges.

A user (that is, someone who needs to view or manage system operations) can have one of these roles:

- Administrator
- Security Administrator
- Manager
- Monitor
- Replication
- Local replication
- Replication/recovery

Note

To issue security CLI commands for user account setup, modification, or removal you must have administrator privileges. Global administrators can manage both global and local user accounts. Local administrators can manage only local user accounts. Table 4 provides an explanation of operations that can be performed by role and scope.

You can assign a user a role globally (the user has the role across all systems in the domain) or locally (the user has the role on a specific system only). Each global username must be unique in the domain; each local username must be unique within the local management server.

Note

Username and passwords can be 1 to 32 characters, including letters (case sensitive), numbers, underscores, non-numeric, non-alpha characters, spaces, and must start with a letter.

[Operations that users with different roles can perform on page 240](#) defines the operations a user can perform by role and scope.

Table 12 Operations that users with different roles can perform

User role	Can view	Can add, modify, or delete
Administrator	All domain and system settings, and global and local accounts	All domain and system settings, and global and local accounts (cannot delete the last global administrator account)
Currently, the following security administrator roles are supported only on CX and CX3 series systems running VNX OE for Block version 02.26.xxx.yyy.zzz or higher.		
Security Administrator	All security-related features, global and local accounts	All security features. Cannot see or manage system features
Manager	All system settings in domain	All system settings in domain
Monitor	All system settings in domain	Nothing
Local Replication	All system settings in domain	Some operations of MirrorView, SAN Copy, and SnapView (see Actions permission on page 240)
Replication	All system settings in domain	Some operations of MirrorView, SAN Copy, and SnapView (see Actions permission on page 240)
Replication and recovery	All system settings in domain	Some operations of MirrorView, SAN Copy, and SnapView (see Actions permission on page 240)

Table 13 Actions permission

Command	Local replication	Replication	Replication/ Recovery
SnapView			
Start a (consistent) snap session	Yes	Yes	Yes
Stop a (consistent) snap session	Yes	Yes	Yes
Activate a session to a snapshot LUN	Yes	Yes	Yes
Deactivate a session from a snapshot LUN	Yes	Yes	Yes

Table 13 Actions permission (continued)

Command	Local replication	Replication	Replication/ Recovery
Synchronize a clone	Yes	Yes	Yes
Fracture a clone	Yes	Yes	Yes
Rollback a snap session	No	No	Yes
Reverse synchronize a clone	No	No	Yes
MirrorView			
Synchronize a mirror/consistency group	No	Yes	Yes
Fracture a mirror/consistency group	No	No	Yes
Control the update parameters of an asynchronous mirror	No	Yes	Yes
Modify the update frequency of an asynchronous mirror	No	Yes	Yes
Throttle a mirror/consistency group	No	Yes	Yes
Promote a synchronous or asynchronous secondary mirror/consistency group	No	No	Yes
SAN Copy			
Start a session	No	Yes	Yes
Stop a session	No	Yes	Yes
Pause a session	No	Yes	Yes
Resume a session	No	Yes	Yes
Mark a session	No	Yes	Yes
Unmark a session	No	Yes	Yes
Verify a session	No	Yes	Yes
Throttle a session	No	Yes	Yes

About LDAP and Active Directory

Unisphere supports the Lightweight Directory Access Protocol (LDAP). Since Microsoft Active Directory is based on LDAP, Active Directory (AD) is supported as well. LDAP uses an active and operational LDAP environment, which allows systems to share LDAP user accounts and username/password authentication.

The administrator and security administrator can then manage the external servers. Users can use their LDAP credentials to log in and perform CLI operations. You must map groups and/or users from the LDAP-compliant directory service to their roles (Administrator and Monitor) using the `security -ldap` commands.

Use these commands to specify the LDAP or AD server you intend to use in your domain, and to map roles to users and user groups in the LDAP server directory.

This feature also supports advanced functions, such as specifying attributes for which user IDs and user/group common names are appended in the LDAP/AD server, defining user and group object classes, and mapping roles to additional servers.

Consult with your site LDAP administrator to determine the correct parameters.

[naviseccli domain and security commands on page 239](#) provides the list of `security -ldap` commands used to configure the LDAP settings.

Domain and security commands overview

The commands in this section include the domain and security commands. The following table identifies the commands.

Table 14 naviseccli domain and security commands

Command	Description
Domain commands	
<code>domain -add</code>	Adds one or more systems to a domain.
<code>domain -list</code>	Displays all systems in a domain.
<code>domain -remove</code>	Removes a system from a domain.
<code>domain -setmaster</code>	Sets a system as the master system in a domain.
<code>domain -uninitialize</code>	Destroys the domain and security configuration of the system.
<code>ntp -list</code>	Lists the NTP configuration settings in a domain.
<code>ntp -set</code>	Modifies the NTP configuration in a domain.
Security commands	
<code>security -adduser</code>	Creates a user account.
<code>security -changeuserinfo</code>	Changes the password and/or role of a user account.
<code>security -list</code>	Displays user accounts.
<code>security -listrole</code>	Lists the user's security role.
<code>security -rmuser</code>	Deletes a user account.
Security LDAP commands	
<code>security -ldap -addServer</code>	Creates a new LDAP directory server.
<code>security -ldap -modifyserver</code>	Modifies the external server login configuration.
<code>security -ldap -removeserver</code>	Removes a server.
<code>security -ldap -listserver</code>	Lists the LDAP external server information.
<code>security -ldap -addrolemapping</code>	Creates a new role mapping.
<code>security -ldap -modifyrolemapping</code>	Modifies the role mapping.
<code>security -ldap -removerolemapping</code>	Removes the role mapping entity.
<code>security -ldap -listrolemapping</code>	Lists the role mapping entity.

Table 14 naviseccli domain and security commands (continued)

Command	Description
<code>security -ldap -synchronize</code>	Synchronizes the accounts.
Security Certificate commands	
<code>security -certificate -add</code>	Imports the certificate.
<code>security -certificate -remove</code>	Removes the specified certificate.
<code>security -certificate -list</code>	Lists all certificates.
<code>security -certificate -alert -BeforeExpiration</code>	Indicates that an alert should be generated prior to the expiration of the SSL/TLS certificate.
<code>security -pkcs12upload</code>	Uploads a PKCS #12 file which contains a certificate and private key to the storage system.
<code>security -SPcertificate -generate</code>	Generates a new self-signed certificate and replaces the one currently in use.
Security FIPS commands	
<code>security -fipsmode -set</code>	Sets the FIPS mode of Unisphere.
<code>security -fipsmode -get</code>	Gets the FIPS mode of Unisphere.
Other Security commands	
<code>security -setport</code>	Sets the http/https ports that Unisphere uses.
<code>security -trustedclient -enable</code>	Enables or disables IP Filter for all storage systems supporting filtering.
<code>security -trustedclient -add</code>	Adds an IP Filter to the storage system.
<code>security -trustedclient -list</code>	Lists the current IP Filters for the storage system.
<code>security -trustedclient -re</code>	Removes IP Filters from the storage system.

domain -add

Adds system to a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli domain` command with the `-add` function, adds the system you specify to the system domain. You can add only one system to the domain one time. When you add a system to the domain, you specify the system's IP address.

See `domain -setmaster` if you need to establish a domain. You then can add system to the domain using `domain -add`.

The `-olduser`, `-oldpassword`, and `-oldscope` are optional switches. These switches must be provided all together or none of them must be provided. If none of them are provided, the user account, `sysadmin` with global scope is used as a default user to operate on the system.

Only one IP address is supported. If multiple IP addresses are provided, displays the error message,

```
Invalid Command. Too many parameters Please also make sure only ONE ip address is provided.
```

If the user uses `domain -add` to add File or Unified systems into a domain without System account, the command fails and displays the error message,

```
Domain add operation failed. The local domain does not have a system account, which is required when adding a file or unified system into a domain. Create a system administrator account in order to continue with adding this system into the domain.
```

Any File or Unified system is allowed to be added into another domain, only if it is in a domain. Otherwise, displays the error message,

```
Domain add operation failed.Security is not initialized. Security must be initialized before any domain operations can be performed in this system. Create a global administrator to initialize security. 10.244.212.164.
```

SYNTAX

```
domain -add IPAddr [[-olduser olduser][--oldpassword oldpassword][--oldscope oldscope]] [-o]
```

OPTIONS

IPAddr

The IP address of the system you want to add.

Note

You must identify the system by the system's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

`-olduser olduser`

Specifies a valid username in the old domain.

`-oldpassword oldpassword`

Specifies the matching password for the `-olduser`.

`-oldscope oldscope`

Specifies the valid scope for the `-olduser`.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE #1

The following example adds the specified systems, to the system domain:

```
naviseccli -address ssl_spa domain -add 111.222.33.44 111.222.33.45
```

```
WARNING: You are about to add following node(s) to the domain.
111.222.33.44
111.222.33.45
Proceed? (y/n) y
```

EXAMPLE #2

```
naviseccli -h 10.14.22.33 -user a -password a -scope global domain
-add 10.14.33.44 -olduser b -oldpassword b -scope global
```

EXAMPLE #3

If any of the switches, `-olduser`, `-oldpassword`, or `-oldscope` are missing, the following error messages are displayed accordingly:

```
C:\>naviseccli -h 10.244.212.164 domain -add 10.244.211.204 -oldscope
global
```

```
Invalid command line parameters. Please provide the scope for the
remote system: -oldscope
```

```
C:\>naviseccli -h 10.244.212.164 domain -add 10.244.211.204 -olduser
global
```

```
Invalid command line parameters. Please provide the user name for
the remote system: -olduser
```

```
C:\>naviseccli -h 10.244.212.164 domain -add 10.244.211.204 -
oldpassword global
```

```
Invalid command line parameters. Please provide the password for the
remote system: -oldpassword
```

OUTPUT

None if the command succeeds; status or error information if it fails.

domain -list

Displays all systems in a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli domain` command with the `-list` function, lists the IP address, name, port and secure port information for all systems in a system domain. If you want to view only information about the master system in the domain, you can use the optional `-master` switch.

Note

You can have only one master system in a domain (see [domain -setmaster on page 248](#)).

SYNTAX

```
domain -list [-master]
```

OPTIONS**-master**

Displays only information about the master system in the domain.

EXAMPLE

The following example lists the systems in the system domain:

```
naviseccli -address 111.222.33.44 domain -list
```

OUTPUT

```
Node:          APM00011111111
IP Address:    111.222.33.55
(Master)
```

```
Name:          CX300I_33_55
Port:          80
Secure Port:   443
IP Address:    111.222.33.44
Name:          CX300I_33_44
Port:          80
Secure Port:   443
```

domain -remove

Removes a system from a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. The software package VNX-Block-Operating-Environment must be committed before this operation is allowed.

DESCRIPTION

The `naviseccli domain` command with the `-remove` function removes the system you specify from a system domain. When you remove a system from a domain, you specify the system's IP address.

If you want to move a system to a different domain, use `domain -remove` to remove the system from the current domain, then `domain -add` to add the system to the other domain.

Note

If you remove the master system from the domain, you destroy the domain. The CLI warns you before you can complete the operation. If the system you want to move is the master system in the domain, and you want to retain the domain, then you should change the master system to another system in the domain before moving the old master (see [domain -setmaster on page 248](#)).

SYNTAX

```
domain -remove IPAddr[-passwordpassword] [-o]
```

OPTIONS

IPAddr

The IP address of the system you want to remove.

Note

You must identify the system by the system's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

`-password password`

Overrides the default password for the global administrator.

NOTICE

You must override the default password. If you don't override it, the system accepts the default password that may pose a high security risk.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example removes the specified system, from the system domain:

```
naviseccli -address ssl_spa domain -remove 111.222.33.44
```

```
WARNING: You are about to remove the following node
from the domain: 111.222.33.44
```

```
Proceed? (y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

domain -uninitialize

Destroys the domain and security configuration of the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli domain` command with the `-uninitialize` function destroys the domain and security configuration of the system you specify. When you destroy the domain and the security configuration of a system, you specify the system's IP address.

Note

You can execute this command only on VNX Block-Only systems with a domain containing only one VNX.

SYNTAX

```
domain -uninitialize IPAddr [-o]
```

OPTIONS*IPAddr*

The IP address of the system whose domain and security configuration you need to destroy.

-o

Executes the command without prompting for confirmation.

OUTPUT

```
E:\chenj20\temp\V1.1
```

```
naviseccli -h 10.244.211.132 domain
```

```
Invalid Command. Too few parameters
usage: domain
-add IPADDRESS [-olduser olduser] [-oldpassword oldpassword]
[-oldscope oldscope]
<-o>
-remove IPADDRESS [-password password] <-o>
-uninitialize IPADDRESS <-o>
-setmaster IPADDRESS <-o>
-list [-master]
```

```
naviseccli -h 10.244.211.132 domain -uninitialize 10.244.211.132
```

```
WARNING: You are about to remove the domain and security
configuration. Proceed? (y/n) y
```

domain -setmaster

Sets a system as the master system in a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli domain` command with the `-setmaster` function, establishes the system you specify as the master system in a domain. The master system holds the master copy of all global account information in the domain. This information is copied to all member systems in the domain. You can have only one master system in a domain. You can use `domain -setmaster` to establish a domain or to change the master system in an existing domain.

If you are using `domain -setmaster` to set up a domain, specify the IP address of a new or existing system that is not part of an existing domain. This system becomes the master and initializes the domain. You can then use the `domain -add` command to add more systems to the domain.

When you use `domain -setmaster` to change the master system in an existing domain, the existing master system is automatically demoted, and the system you specify becomes the new master system.

SYNTAX

```
domain -setmaster IPAddr[-o]
```

OPTIONS

IPAddr

The IP address of the system you want to set as master.

Note

You must identify the system by the system's IP address. Specify the IP address in the form b.c.d.e where b, c, d and e range from 0 through 255, for example, 111.222.33.44.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example sets the specified system as the master system in a domain:

```
naviseccli -address ssl_spa domain -setmaster 111.222.33.44
```

```
WARNING: You are about to set the following node as
the master of the domain: 111.222.33.44
```

```
Proceed? (y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

ntp -list

Lists the NTP configuration settings in effect in a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Only a security administrator or an administrator has the access rights.

You will receive an error message if access is denied or an invalid parameter is set.

DESCRIPTION

The `naviseccli ntp` command with the `-list` function, lists the Network Time Protocol (NTP) configuration settings in effect in a domain.

SYNTAX

```
ntp -list [-control|-servers|-all
```

OPTIONS

`-control`

Displays start/stop and interval information.

`-servers`

Displays configured server addresses and any associated authentication information.

`-all`

Displays all information.

EXAMPLE

The following example lists all the information:

```
naviseccli -user a -password a -scope 0 -address 10.5.1.207 ntp -list
-all
```

OUTPUT

```
Output: -all
start: YES
interval: 720 minutes
address: 10.5.1.207 128.221.142.13
serverkey: 0
keyvalue: ""

XML Output:

-list -all

<?xml version="1.0" encoding="utf-8" ?>
<CIM CIMVERSION="2.0" DTDVERSION="2.0"><MESSAGE ID
="877"
PROTOCOLVERSION="1.0"><SIMPLERSP><METHODRESPONSE
NAME="ExecuteClientRequest"><RETURNVALUE
TYPE="Navi_Error">
<VALUE.NAMEDINSTANCE>
<INSTANCENAME CLASSNAME="Navi_Error">
</INSTANCENAME>
<INSTANCE CLASSNAME="Navi_Error">
<PROPERTY NAME="errorCode"
TYPE="uint32"><VALUE>0</VALUE>
</PROPERTY>
<PROPERTY NAME="success"
TYPE="boolean"><VALUE>>true</VALUE>
</PROPERTY>
<PROPERTY NAME="where"
TYPE="string"><VALUE>NTPProvider</VALUE>
</PROPERTY>
<PROPERTY NAME="why" TYPE="string"><VALUE>Operation
successfully completed. </VALUE>
</PROPERTY>
</INSTANCE>
</VALUE.NAMEDINSTANCE>
</RETURNVALUE><PARAMVALUE NAME="start:"
```

```

TYPE="string"><VALUE>yes</VALUE>
</PARAMVALUE>
<PARAMVALUE NAME="interval:"
TYPE="string"><VALUE>30</VALUE>
</PARAMVALUE>
<PARAMVALUE NAME="address:"
TYPE="string"><VALUE>10.5.4.211 128.222.132.13</VALUE>
</PARAMVALUE>
<PARAMVALUE NAME="serverkey:" TYPE="string"><VALUE>1
0</VALUE>
</PARAMVALUE>
<PARAMVALUE NAME="keyvalue:"
TYPE="string"><VALUE>"1234567812345678"
&quot;&quot;</VALUE>
</PARAMVALUE>
</METHODRESPONSE></SIMPLERSP></MESSAGE></CIM>

```

ntp -set

Modifies the NTP configuration in a domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Only a security administrator or an administrator has access rights.

You will receive an error message if the access is denied or an invalid parameter is set.

DESCRIPTION

The `naviseccli ntp` command with the `-set` function modifies the NTP configuration in a domain.

SYNTAX

```
ntp -set [-start yes|no] [-interval n -servers addr1 [addr2...]
-serverkey sk1 [sk2...] -keyvalue kval1 [kval2...]]
```

OPTIONS

`-start yes|no`

Yes starts NTP; no stops it. To start NTP, servers must have been defined and have been displayed previously in a list command or while running this command.

`-interval n`

Specifies the synchronization interval in minutes ranging from 30 to 43200 (30 days).

`-servers addr1 [addr2...]`

Specifies the list of IPv4 NTP server addresses. At least one address must be specified.

`-serverkey sk1 [sk2...]`

Specifies the list of the integer values (1-65534 inclusive). This option is required if server authentication is desired. An ordered list is in the same order as `-servers`. A value of 0 indicates no key will be used. In that case use a keyvalue of "". If `serverkey` is used, every address entry must have a corresponding server key.

`-keyvalue kval1 < kval2...`

Specifies the list of key values. This option is required if `serverkey` is specified. If server authentication is not desired, `keyvalue` is not required. Enclose key values in quotation marks. An empty value is simply a pair of quotation marks (for `serverkey` value = 0). An ordered list is in the same order as `serverkey` above. Every server key must have a key value. Valid key values are printable ASCII characters excluding the space character and the # character. If specified, the length of each key must be between 1 and 16 characters (empty value is described above).

EXAMPLE #1

Specifies all items for `ss1_spa` including server authentication.

```
naviseccli -address ss1_spa -set -start yes -interval
45 -servers 10.5.1.207 128.222.132.13 -serverkey
1 5 -keyvalue "1234567812345678" "abcdefghabcdefgh"
```

EXAMPLE #2

Specifies all items for `ss1_spa` including server authentication but the key for `addr2` is not specified.

```
naviseccli -address ss1_spa -set -start yes -interval
45 -servers 10.5.1.207 128.222.132.13 -serverkey
5 0 -keyvalue "1234567812345678" ""
```

EXAMPLE #3

Specifies items for `ss1_spa` choosing not to use server authentication.

```
naviseccli -address ss1_spa -set -start yes -interval
45 -servers 10.5.1.207
```

EXAMPLE #4

Stops (disables) NTP operations for `ss1_spa`.

Note

Does not affect synchronization interval or servers settings.

```
naviseccli -address ss1_spa -set -start no
```

EXAMPLE #5

Specifies the interval and server information for `ss1_spa`.

Note

Does not affect start/stop settings.

```
naviseccli -address ss1_spa -set -interval 45 -servers
10.5.1.207 128.222.132.13 -serverkey 5 0 -keyvalue
"1234567812345678" ""
```

OUTPUT

Varies depending on the switch you use.

security -adduser

Creates a user account.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

The software package VNX-Block-Operating-Environment must be committed before this operation is allowed.

DESCRIPTION

The `naviseccli security` command with the `-adduser` function, adds a user account to the system you specify. If you create a global account, it replicates to all systems in the domain. When you add a user account, you specify the username, password, scope and role.

System account is a special global administrator account which is used to issue CLI command from File to Block. System account needs to be created immediately after the upgrade to R31

The `-type` option differentiates the user account and system account. If `-type` is not specified, the system creates the account of `User` type, by default.

Note

When you create the initial global administrator user account, you do not need to provide a username and password at login.

SYNTAX

```
security -adduser -user username -password password -scope
global|local -role administrator|storageadmin|operator|
securityadministrator|dataprotection|localdataprotection|
datarecovery|sanadmin|networkadmin|nasadmin|vmadmin [-type
user|system] [-o]
```

OPTIONS

`-user username`

Specifies the username you want to create for the user account.

`-password password`

Specifies the password you want to create for the user account.

`-scope global|local|LDAP`

Specifies the scope (global, local, or LDAP) you want to apply to the user account.

`-role administrator|storageadmin|operator|securityadministrator|dataprotection |
localdataprotection|datarecovery|sanadmin|networkadmin|nasadmin|vmadmin]`

Specifies the role, administrator, storageadmin, operator, securityadministrator,
dataprotection, localdataprotection, datarecovery, sanadmin, networkadmin,
nasadmin, or vmadmin you want to apply to the user account.

`-type user|system`

If you specify user, a user account is created. If you specify system, a system account
is created. You cannot create more than one system account in a domain and a non-
global or a non-administrator system account.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example creates the initial global administrator account on the system. You are then prompted to create a domain for the system you specify:

```
naviseccli -address ssl_spa security -adduser -user a -password
b -scope global -role administrator
```

```
Global security is not initialized. It is highly recommended
that you initialize global security.
```

```
Do you still want to continue? y
```

```
WARNING: You are about to add user: a
```

```
Proceed?(y/n) y
```

```
This system is not in a domain. It is highly
recommended you create a new domain for this system.
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -changeuserinfo

Changes the password and/or role of a user account.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security` command with the `-changeuserinfo` function, modifies the password and/or role of a user account. When you modify a user account, you specify the username and scope, then designate a new password and/or role.

Note

You can change the password in a user account without providing the user's current password.

Note

If a user attempts to change the role or the scope of an account with the System designation, the command fails.

SYNTAX

```
security -changeuserinfo -user username -scope global|local
[-newpassword password] [-newrole administrator|manager|
monitor]
[-o]
```

OPTIONS

`user` *username*

Specifies the username of the user account you want to modify.

`-scope` *global|local*

Specifies the scope, global or local, of the user account you want to modify. A global account has access to all systems in a domain.

`-newpassword` *password*

Specifies the new password you want to create for the user account.

`-newrole` *administrator|manager|monitor*

Specifies the new role, administrator, manager, or monitor, you want to apply to the user account.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example modifies the role of an existing user account:

```
naviseccli -address ssl_spa security -changeuserinfo -user
b -scope local -newrole manager
```

```
WARNING: You are about to change user: b (local)
```

```
Proceed? (y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -list

Displays user accounts.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security` command with the `-list` function, lists the username, scope, role of user accounts, and type of account.

Note

The `security -list` command is supported on all Unisphere systems.

SYNTAX

```
security -list [-user username] [-scope global|local] [-role
administrator|storageadmin|operator|securityadministrator|
dataprotection |localdataprotection|datarecovery|sanadmin|
networkadmin|nasadmin|vmadmin] [-type]
```

OPTIONS

`-user username`

Identifies a specific username for which you want to display information.

`-scope global|local`

Identifies users with a specific scope for which you want to display information.

`-role`

Identifies users with a specific role for which you want to display information.

`-type`

Displays the type of account, either a user account or a system account.

EXAMPLE

The following example lists users with the specified scope and role:

```
naviseccli -address ssl_spa security -list -role
monitor -scope global
```

OUTPUT

```
Username: c
Role:     monitor
Scope:   global
```

security -listrole

Lists the user's security role.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. If no user account matches the query, an access denied error will be displayed.

DESCRIPTION

The `naviseccli security` command with the `-listrole` function, lists the user's security role.

SYNTAX

```
security -listrole
```

EXAMPLE

The following example lists the user's security role:

```
naviseccli -h 10.14.83.44 security -listrole
```

OUTPUT

```
Role: Administrator
```

security -rmuser

Deletes a user account.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security` command with the `-rmuser` function, removes the user account you specify. When you remove a user account, you specify the username and scope.

Note

You must specify a scope for the user you want to remove, since you can establish two user accounts with the same username and varying scopes, one global and one local.

SYNTAX

```
security -rmuser -user username -scope global|local [-o]
```

OPTIONS

`-user username`

Identifies the username of the account you want to remove.

`-scope global|local`

Specifies the scope, global or local, of the account you want to remove.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The following example removes the specified user account:

```
naviseccli -address ssl_spa security -rmuser -user b -scope
local
```

```
WARNING: You are about to remove user: b (local)
```

```
Proceed? (y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -addserver

Creates a new LDAP directory server.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-addserver` subcommand creates a new external server login configuration. Only a Unisphere security administrator or administrator can connect your system to an LDAP or Active Directory (AD) service. After defining the server connection settings, and mapping the user's roles, users can log in to the system with their LDAP username/password. Only two service connections can be configured at a time.

System security must be initialized and a domain master must be selected before configuring LDAP connection settings. Consult with the LDAP site administrator to determine the correct parameters.

SYNTAX

```
security -ldap -addserver IPaddress -portnumber portnumber
-servertype LDAP|AD -protocol LDAP|LDAPS -binddn binddn
-bindpassword password -usersearchpath usersearchpath
[-groupsearchpath groupsearchpath] [-useridattribute attribute]
[-usernameattribute nameattribute] [-groupnameattribute
groupattribute] [-groupmemberattribute memberattribute]
[-userobjectclass objectclass] [-groupobjectclass groupclass]
[-certpathToCertificatefile] [-o]
```

OPTIONS

`IPaddress`

Specifies the network IP address (IPv4) of the LDAP/AD server with the information required to authenticate users.

`-portnumber portnumber`

Specifies the port number that the server will open to allow external programs to access it. If the protocol is LDAP, the default port number is 389. If the protocol is LDAPS, the default port number is 636.

`-servertype LDAP|AD`

Specifies whether the server is a LDAP directory or AD server. The default is LDAP.

`-protocol LDAP|LDAPS`

Lets you specify either LDAPS or LDAP protocol.

`-binddn binddn`

Specifies the login name for the LDAP/AD server; it cannot exceed 512 characters. This field allows an LDAP-enabled application to access the targeted server and search for a user ID/password to validate. It should be in "cn=, ou=, dc=, dc=" format.

`-bindpassword password`

Credentials used to authenticate the bind DN. If you specify the length of other options such as `-binddn`, the length of `bindpassword` should also be specified here. The `bindpassword` cannot exceed 512 characters.

`-usersearchpath usersearchpath`

Sets the path within the LDAP user entry database where a search for the username/password begins. It cannot exceed 512 characters. It should be in "ou=, dc=, dc=" format.

`-groupsearchpath groupsearchpath`

Sets the path within the LDAP user-group entry database where a search for the group name/password begins. It cannot exceed 512 characters. It should be in "ou=, dc=, dc=" format. The option defaults to the User Search Path if you do not specify the path.

`-useridattribute useridattribute`

Specifies the attribute to which the user ID will be appended in the LDAP/AD servers. The directory hierarchy will be searched using this attribute/userid pair. The default

for an LDAP directory is uid and the default for an AD is sAMAccountName. It cannot exceed 128 characters.

`-usernameattribute` *nameattribute*

Specifies the attribute to which the user's common name (cn) will be appended in the servers. The default is cn and it cannot exceed 128 characters.

`-groupnameattribute` *groupattribute*

Specifies the attribute to which the user group's common name will be appended in the servers. This is stored as an attribute of an entry within the Group Search Path of the directory. The default is cn and it cannot exceed 128 characters.

`-groupmemberattribute` *groupattribute*

Acts as a search filter for the different attribute types to identify the different groups of members. If you do not specify this field, then a search will be performed using only the Group Name Attribute parameter. The default value for an LDAP directory is uniqueMembers and the default for an AD is member. It cannot exceed 128 characters.

`-userobjectclass` *objectclass*

Defines the required and optional attributes so that the user entry can act as a search filter in a situation where a user has multiple entries in a server. If you do not specify this field, then a search will be performed using only the User ID Attribute parameter. It cannot exceed 128 characters.

`-groupobjectclass` *groupclass*

Defines the required and optional attributes so that the group entry can act as a search filter in a situation where a group has multiple entries in a server. If you do not specify this field, then a search will be performed using only the Group Name Attribute parameter. The default for an LDAP directory is groupOfUniqueNames and the default for an AD is group. It cannot exceed 128 characters.

`-cert` *pathToCertificatefile*

Specifies the full pathname of the trusted certificate file to be uploaded to the certificate store.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -h 10.32.123.229 -user a -password a -scope 0 security
-ldap -addserver 10.5.4.111 -portnumber 389 -servertype LDAP
-protocol LDAP -binddn cn=Manager,dc=ipv4,dc=com
-bindpassword ipv4int -usersearchpath ou=Users,dc=ipv4,dc=com
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -modifyserver

Modifies the external server login configuration.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-modifyserver` subcommand modifies the settings that connect your system to an LDAP or AD service. Consult with the LDAP/AD site administrator to determine the correct parameters. Only the security administrator or administrator can execute this command.

SYNTAX

```
security -ldap -modifyserver IPAddress -bindpassword password
[-ldapservers IPAddress] [-portnumber portnumber] [-servertype
LDAP| AD] [-protocol LDAP| LDAPS] [-binddn binddn]
[-newbindpassword password] [-usersearchpath usersearchpath]
[-groupsearchpath groupsearchpath] [-useridattribute
useridattribute]
[-usernameattribute usernameattribute] [-groupnameattribute
groupnameattribute] [-groupmemberattribute
groupmemberattribute] [-userobjectclass userobjectclass]
[-groupobjectclass groupclass] [-o]
```

OPTIONS

IPAddress

Specifies the network IP address (IPv4) of the LDAP/AD server with the information required to authenticate users.

-bindpassword password

Credentials used to authenticate the bind DN. If you specify the length of other options such as *-binddn*, the length of *bindpassword* should also be specified here. The *bindpassword* cannot exceed 512 characters.

-ldapservers IPAddress

Specifies the network IP address (IPv4) of the LDAP server.

-portnumber portnumber

Lets you specify the port number that the server will open to allow external programs to access it. If the protocol is LDAP, the default port number is 389. If the protocol is LDAPS, the default port number is 636.

-servertype LDAP|AD

Specifies whether the server is an LDAP directory or AD server. The default is LDAP.

-protocol LDAP|LDAPS

Lets you specify either LDAPS or LDAP protocol.

-binddn binddn

Specifies the login name for the LDAP/AD server; it cannot exceed 512 characters. This field allows an LDAP-enabled application to access the targeted server and search for a *userid/password* to validate. It should be in "cn=, ou=, dc=, dc=" format.

-newbindpassword password

Specifies the password that corresponds to the bind DN; and it cannot exceed 512 characters.

-usersearchpath usersearchpath

Sets the path within the LDAP user entry database where a search for the *username/password* begins. It cannot exceed 512 characters. It should be in "ou=, dc=, dc=" format.

-groupsearchpath groupsearchpath

Sets the path within the LDAP user-group entry database where a search for the *group name/password* begins. It cannot exceed 512 characters. It should be in "ou=, dc=, dc=" format. The option defaults to the User Search Path if you do not specify the path.

-useridattribute useridattribute

Specifies the attribute to which the user ID will be appended in the LDAP/AD servers. The directory database is searched using this attribute/*userid* pair. The default for an LDAP directory is *uid* and the default for an AD is *sAMAccountName*. It cannot exceed 128 characters.

-usernameattribute usernameattribute

Specifies the attribute to which the user's common name (cn) will be appended in the servers. The default is *cn*. It cannot exceed 128 characters.

- groupnameattribute *groupnameattribute*
Specifies the attribute to which the user group's common name will be appended in the servers. This is stored as an attribute of an entry within the Group Search Path of the directory. The default is cn. It cannot exceed 128 characters.
- groupmemberattribute *groupmemberattribute*
Acts as a search filter for the different attribute types to identify the different groups of members. If you do not specify this field, then a search will be performed using only the Group Name Attribute parameter. The default value for an LDAP directory is uniqueMembers and the default for an AD is member. It cannot exceed 128 characters.
- userobjectclass *userobjectclass*
Defines the required and optional attributes so that the user entry can act as a search filter in a situation where a user has multiple entries in a server. If you do not specify this field, then a search will be performed using only the User ID Attribute parameter. It cannot exceed 128 characters.
- groupobjectclass *groupclass*
Defines the required and optional attributes so that the group entry can act as a search filter in a situation where a group has multiple entries in a server. If you do not specify this field, then a search will be performed using only the Group Name Attribute parameter. The default for an LDAP directory is groupOfUniqueNames, and the default for an AD is group. It cannot exceed 128 characters.
- o
Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -h 10.32.123.229 -user a -password a -scope 0 security
-ldap -modifyserver 10.5.4.250 -bindpassword crosstor -servertype
LDAP
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -removeserver

Removes a server.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-removeserver` subcommand deletes an LDAP external server login configuration. It also destroys all role mapping information related to the specified server. Only the security administrator or administrator can execute the command.

SYNTAX

```
security -ldap -removeserver IPaddress [-o]
```

OPTIONS

IPaddress

Deletes the specified network IP address (IPv4) of the LDAP or the AD server.

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
security -ldap -removeserver 10.5.4.250
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -listserver

Lists the LDAP external server information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap -listserver` command lists the existing service connections, that the IP address of the LDAP server identifies. You can use the optional switches to list a particular server. Only the security administrator or administrator can execute the command.

SYNTAX

```
security -ldap -listserver [-primaryserver|-alternateserver]
```

OPTIONS

`-primaryserver`

Displays the primary server information. The connection, which is set up first, is automatically designated as the primary and is contacted first to authenticate a user or user request.

`-alternateserver`

Displays the alternate server information. The only way to change the alternate server to primary is to delete the primary server.

EXAMPLE**Security -ldap -listserver**

```
Primary Server Name: 10.5.4.111
Primary Server Port Number: 389
Primary Server Protocol: ldap
Primary Server ServerType: AD
Primary Server BindDn:
cn=administrator,cn=users,dc=NavilabTEST,dc=us,dc=dg,dc=
com
Primary Server UserSearchPath:
dc=NavilabTEST,dc=us,dc=dg,dc=com
Primary Server GroupSearchPath: groupsearchpath
Primary Server UserIDAttribute: testuid
Primary Server UserNameAttribute: usernameattribute
Primary Server GroupNameAttribute: groupnameattribute
Primary Server GroupMemberAttribute:
groupmemberattribute
Primary Server UserObjectClass(optional):
userobjectclass
Primary Server GroupObjectClass: groupclass

Alternate Server Name: 10.14.46.100
Alternate Server Port Number: 389
Alternate Server Protocol: ldap
Alternate Server ServerType: LDAP
Alternate Server BindDn: cn=root,dc=corp,dc=emc,dc=com
Alternate Server UserSearchPath:
ou=Group1,dc=corp,dc=emc,dc=com
Alternate Server GroupSearchPath: groupsearchpath
Alternate Server UserIDAttribute: useridattribute
Alternate Server UserNameAttribute: usernameattribute
Alternate Server GroupNameAttribute: groupnameattribute
Alternate Server GroupMemberAttribute:
groupmemberattribute
```

```
Alternate Server UserObjectClass(optional):
userobjectclass
Alternate Server GroupObjectClass: groupclass
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -addrolemapping

Creates a new role mapping.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-addrolemapping` subcommand creates a new role mapping entity for the corresponding external server. You must initialize system security and select a domain master before configuring LDAP connection settings. Consult with the LDAP administrator for your site to determine the LDAP/AD names of users and groups to which you want to map Unisphere roles. Once you have defined the server connection settings, and mapped the user's roles, users can log in to the system with their LDAP username/password. The roles are administrator, manager, monitor, security administrator, replication, local replication, and replication/recovery.

SYNTAX

```
security -ldap -addrolemapping IPaddress -name name -type
group|user -role role [-o]
```

OPTIONS

IPaddress

Specifies the network IP address (IPv4) of the LDAP or AD server with the information required to authenticate users.

-name *name*

Specifies the name of the role that maps to Unisphere.

-type group|user

Specifies the type of role mapping. The group type of role mapping applies to a group of people. The user type of role mapping applies to only one individual.

-role *role*

Specifies the user/group role that maps to Unisphere.

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
security -ldap -addrolemapping 10.5.4.250 -name ldap4ipv4 -type
user -role administrator
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -modifyrolemapping

Modifies the role mapping.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-modifyrolemapping` subcommand modifies the user or group role mappings to an LDAP or AD service.

SYNTAX

```
security -ldap -modifyrolemapping IPaddress -name name -type
group|user -role role [-o]
```

OPTIONS

`IPaddress`

Specifies the network IP address (IPv4) of the LDAP or the AD server.

`-name name`

Specifies the name of the role that maps to Unisphere.

`-type group|user`

Specifies the type of the role mapping to be modified. The group type of role mapping applies to a group of people. The user type of role mapping applies to only one individual.

`-role role`

Specifies the user/group role that maps to Unisphere.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
security -ldap -modifyrolemapping 10.5.4.250 -name test -type user
-role monitor
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -removerolemapping

Removes the role mapping entity.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-removerolemapping` subcommand removes the role mapping entity from Unisphere for the corresponding external server. Only the Unisphere administrator can remove user or group role mappings to an LDAP or AD service.

You can delete all role mappings using the `-all` switch. To delete a particular role mapping, specify the `-name` and `-type` switch.

SYNTAX

```
security -ldap -removerolemapping IPaddress [-name name -type
group|user] | [-all] [-o]
```

OPTIONS

`IPaddress`

Specifies the network IP address (IPv4) of the LDAP or the Active Directory server.

`-name name -type group|user`

Deletes role mapping with a specified role mapping name and type. The group type of role mapping applies to a group of people. The user type of role mapping applies to only one individual.

`-all`

- Deletes all role mappings.
- o Executes the command without prompting for confirmation.

EXAMPLE

```
security -ldap -removerolemapping 10.5.4.250 -name test -type user
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -ldap -listrolemapping

Lists a role mapping entity.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-listrolemapping` subcommand lists the role mapping entity for the corresponding external server. It displays the role mapping as a security administrator, administrator, manager, or monitor. You can use the `-name` switch to list a particular role mapping.

SYNTAX

```
security -ldap -listrolemapping IPaddress [-name name]
```

OPTIONS

IPaddress

Specifies the network IP address (IPv4) of the LDAP or the AD server with the information required to authenticate users.

`-name name`

Displays role mapping with the specified role mapping name.

EXAMPLE

```
security -ldap -listrolemapping 10.5.4.250
```

OUTPUT

```
Name: Test_1
Type: user
Role: administrator
Name: Test_2
Type: group
Role: security administrator
```

security -ldap -synchronize

Synchronizes the accounts.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -ldap` command with the `-synchronize` subcommand manually synchronizes the accounts with the current service.

SYNTAX

```
security -ldap -synchronize [-o]
```

OPTIONS

```
-o
```

Executes the command without prompting for confirmation.

OUTPUT

None if the command succeeds; status or error information if it fails.

security -certificate -add

Imports the certificate.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `security -certificate -add` command imports the certificate from the certificate file. Secure CLI supports BASE64PEM format certificate.

SYNTAX

```
security -certificate -add -file filename [-client -role role]
```

OPTIONS

```
-file filename
```

Specifies the path of the certification on the local file system.

```
-client
```

Specifies that this is a client certificate and not a server certificate.

```
-role role
```

Specifies the role you want to apply to the client certificate.

EXAMPLE

```
naviseccli security -certificate -add -file c:\Node1.cer
```

security -certificate -remove

Removes the specified certificate.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -certificate -remove` command removes the specified certificate in certificate store.

SYNTAX

```
security -certificate -remove [-issuer issuer -serialnumber serial number] [-client] |[-all]
```

OPTIONS

```
-issuer issuer
```

Specifies the issuer of the certificate that will be removed.

```
-serialnumber serial number
```

Specifies the certificate.

```
-all
```

Removes all certificates in the certificate store.

```
-client
```


Specifies that this is a client certificate rather than a server certificate.

EXAMPLE

```
naviseccli security -certificate -remove -issuer "CN=ICAl 1rst
Intermediate CA - Signed by Root, OU=CSP, O=RSA, ST=MA,C=US"
-serialNumber 00e2
```

security -certificate -list

Lists the certificates.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -certificate -list` command lists all certificates in the certificate store. Use the `expire` option to list only the expired certificate.

SYNTAX

```
security -certificate -list
[-client][-showOnly expired]
```

OPTIONS

`-showOnly expired`
Lists only the expired certificates.

`-client`
Lists the client SSL certificates.

EXAMPLE

```
naviseccli security -certificate -list
```

OUTPUT

```
-----
Subject:
CN=TrustedRoot,C=US,ST=MA,L=Hopkinton,EMAIL=rsa@emc.com,OU=CSP,O=RSA
Issuer:
CN=TrustedRoot,C=US,ST=MA,L=Hopkinton,EMAIL=rsa@emc.com,OU=CSP,O=RSA
Serial#:          00d8280b0c863f6d4e
Valid From:       20090407135111Z
Valid To:         20190405135111Z
-----
Subject:
CN=ICAl 1rst Intermediate CA - Signed by Root, OU=CSP,
O=RSA, ST=MA,C=US
Issuer: CN=TrustedRoot,
C=US,ST=MA,L=Hopkinton,EMAIL=rsa@emc.com,OU=CSP,O=RSA
Serial#:          00e2
Valid From:       20090407135113Z
Valid To:         20190405135113Z
-----
Subject: CN=ICAl 2nd Intermediate CA,OU=CSP,O=RSA,ST=MA,C=US
Issuer: CN=ICAl 1rst Intermediate CA - Signed by
Root,OU=CSP,O=RSA,ST=MA,C=US
Serial#:          00e3
Valid From:       20090407135113Z
Valid To:         20190405135113Z
-----
```

```

Subject: CN=node 1 CA - Signed by ICA1, OU=CSP, O=RSA, ST=MA,C=US
Issuer: CN=ICA1 1rst Intermediate CA - Signed by Root, OU=CSP,
O=RSA, ST=MA,C=US
Serial#: 00e2
Valid From: 20090407135113Z
Valid To: 20190405135113Z
-----

```

security -certificate -alertBeforeExpiration

Indicates that an alert should be generated prior to the expiration of the SSL/TLS certificate.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -certificate -alertBeforeExpiration` command indicates that an alert should be generated prior to the expiration of the SSL/TLS certificate.

SYNTAX

```
security -certificate -alertBeforeExpiration DAYS
```

OPTIONS

`-alertBeforeExpiration DAYS`

Sets the number of days to send an alert before the certificate expires. Valid values are from 0 to 366.

EXAMPLE

```
naviseccli security -certificate -alertBeforeExpiration 10
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -pkcs12upload

Uploads a PKCS #12 file which contains a certificate and private key to the storage system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -pkcs12upload` command uploads a PKCS #12 file, which contains a certificate and private key, to the storage system.

SYNTAX

```
security -pkcs12upload -file FILE [-passphrase PASSPHRASE]
```

OPTIONS

`-file FILE`

Specifies the file full name of the certificate to be imported.

`-passphrase PASSPHRASE`

Specifies the export password of the certificate.

EXAMPLE

```
naviseccli security -pkcs12upload c:\cert.p12
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -SPcertificate -generate

Generates a new self-signed certificate and replaces the one currently in use.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -SPcertificate -generate` command generates a new self-signed certificate and replaces the one currently in use.

SYNTAX

```
security -SPcertificate -generate
```

EXAMPLE

```
naviseccli security -SPcertificate -generate
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -fipsmode -set

Sets the FIPS (Federal Information Processing Standards) mode of Unisphere.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -fipsmode -set` command sets the FIPS mode of Unisphere.

SYNTAX

```
security -fipsmode -set 0|1 [-o]
```

OPTIONS

```
-set 0|1
```

Sets the FIPS mode of Unisphere. Valid values are 0 and 1.

0

Sets the mode to Non-FIPS 140-2-compliant mode.

1

Sets the mode to FIPS 140-2-compliant mode.

```
-o
```

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli security -fipsmode -set 1
```

WARNING: You are about to change FIPS mode. Any other existing background management tasks will be interrupted and the system will be unavailable for management while the environment restarts. IO to the system will be unaffected.
Proceed? (y/n)

OUTPUT

None if the command succeeds; status or error information if it fails.

security -fipsmode -get

Gets the FIPS (Federal Information Processing Standards) mode of Unisphere.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -fipsmode -get` command gets the FIPS mode of Unisphere.

SYNTAX

```
security -fipsmode -get
```

EXAMPLE

```
naviseccli security -fipsmode -get
```

OUTPUT

```
FIPS 140-2 compliant: ENABLED
```

security -setport

Sets the http/https ports that Unisphere uses.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -setport` command changes the http/https port numbers.

SYNTAX

```
security -setport -port [80/443] |[2162/2163]
```

OPTIONS

```
-port [80/443] |[2162/2163]
```

Sets the ports that Unisphere uses. Valid port pairs are either 80/443 or 2162/2163. EMC recommends that all systems in the domain use the same management port pair.

EXAMPLE

```
naviseccli security -setport -port 80/443
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -trustedclient -enable

Enables or disables IP Filtering for all storage systems supporting filtering.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -trustedclient` command with the `-enable` subcommand enables or disables IP Filtering for all storage systems supporting filtering.

SYNTAX

```
security -trustedclient -enable yes|no [-o]
```

OPTIONS

`-enable yes|no`

Enables (yes) or disables (no) IP Filtering.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli security -trustedclient -enable yes
```

```
WARNING: You are about to enable IP Filtering which may cut off
access from clients.
Proceed?(y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -trustedclient -add

Adds an IP Filter to the storage system or Unisphere Domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -trustedclient` command with the `-add` subcommand adds an IP Filter to the storage system or Unisphere Domain. Administrators and security administrators can extend remote service tool access to a storage system's management ports by entering the IP addresses of attached, trusted clients.

SYNTAX

```
security -trustedclient -add -ipaddress ADDRESS -applyto
single|all [-o]
```

OPTIONS

`-ipaddress ADDRESS`

Specifies the client's IP Address to allow access for.

`-applyto single|all`

Applies the filter to the connected storage system only (single) or all storage systems in the domain (all).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli security -trustedclient -add -ipaddress 10.238.74.171 -
applyto all
```

```
WARNING: You are about to grant access solely to this trusted client.
Proceed?(y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

security -trustedclient -list

Lists the current IP Filters for the storage system and Unisphere Domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -trustedclient` command with the `-list` subcommand lists the current IP Filters for the storage system and Unisphere Domain.

SYNTAX

```
security -trustedclient -list [-ipaddress ADDRESS][-showall]
```

OPTIONS

`-ipaddress ADDRESS`

Specifies the client's IP Address to list filtering for.

`-showall`

Shows all trusted clients, including locals that only apply to other storage systems.

EXAMPLE

```
naviseccli security -trustedclient -list
```

OUTPUT

```
Current IP Filtering State: enabled
Trusted Client: 10.238.74.171
Applies To: all
```

security -trustedclient -remove

Removes IP Filters from the storage system or Unisphere Domain.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli security -trustedclient` command with the `-remove` subcommand removes IP Filters from the storage system or Unisphere Domain.

SYNTAX

```
security -trustedclient -remove [-ipaddress ADDRESS] -applyto
single|all [-o]
```

OPTIONS

`-ipaddress ADDRESS`

Specifies the client's IP Address to remove filtering for.

`-applyto single|all`

Applies the filter to the connected storage system only (single) or all storage systems in the domain (all).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli security -trustedclient -remove -ipaddress 10.238.74.171 -
applyto all
```

```
WARNING! If you proceed, access will be terminated for the initiating
client.
Proceed?(y/n) y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

Event Monitor commands

This section provides a brief overview of the event monitor and includes a list of commands you can use to configure and manage Event Monitor.

Event Monitoring overview

The Event Monitor is an enterprise tool that supports centralized or distributed monitoring of storage systems in a heterogeneous environment. Event Monitor is part of the Agent, and it is available on many operating systems. Once configured, the Event Monitor runs continuously as a service or daemon, observing the state of all specified storage systems and notifying you when selected events have occurred.

The Event Monitor has portal, template, and monitor commands for portal configuration, template management, and centralized and distributed event monitor management respectively.

[naviseccli Event Monitor commands on page 271](#) lists the event monitor commands.

Table 15 naviseccli Event Monitor commands

Command	Description
-portal -create	Adds a host agent to the portal system.
-portal -list	Lists the portal configurations.
-portal -migrate	Migrates the portals.
-portal -destroy	Deletes the host agent from the portal system.
-template -list	Lists the templates.
-template -import	Imports the template files.
-template -export	Exports the templates.
-template -destroy	Deletes the templates.
-template -findconflict	Finds and resolves template conflicts.
-template -swap	Swaps two templates.
-monitor -addsystem	Adds the storage system to be monitored.
-monitor -removesystem	Removes the storage system from the monitor.
-monitor -listmapping	Lists the storage template mappings.
-monitor -applytemplate	Applies the template to the monitor.

Table 15 naviseccli Event Monitor commands (continued)

Command	Description
<code>-monitor -stoptemplate</code>	Stops applying the template to the monitor.
<code>-monitor -getlog</code>	Displays the Event Monitor log.
<code>-monitor -enableresponse</code>	Enables the response from the specified monitor.
<code>-monitor -disableresponse</code>	Disables the response from the specified monitor.
<code>-monitor -getlogsize</code>	Gets the log size of the monitor.
<code>-monitor -setlogsize</code>	Sets the log size of the monitor.
<code>-monitor -insertttestevent</code>	Inserts an event in the log.
<code>-monitor -responsetest</code>	Creates an artificial event.
<code>-monitor -reloadconfig</code>	Reloads the template file.

-portal -create

Adds a host agent to the portal system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -portal -create` command adds the specified host agent to the specified portal system.

SYNTAX

```
eventmonitor -portal -create -system portalsystemname -server hostagentname
```

OPTIONS

```
-system portalsystemname
    Specifies the portal system.
-server hostagentname
    Specifies the host agent.
```

EXAMPLE

```
eventmonitor -portal -create -system cx380_123_227 -server 10.32.123.205
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-portal -list

Lists the portal configurations.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -portal -list` command lists the portal system and its hosts.

SYNTAX

```
eventmonitor -portal -list
```

OUTPUT

Portal configuration is listed. Nothing is listed if the portal configuration is not available.

-portal -migrate

Migrates the portals.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -portal -migrate` command migrates the portal system. The host to which the portals are connected is migrated automatically.

SYNTAX

```
eventmonitor -portal -migrate -source source portal system name
-dest destination portal system name
```

OPTIONS

`-source` *source portal system name*

Specifies the original portal configuration name.

`-dest` *destination portal system name*

Specifies the destination portal configuration name.

EXAMPLE

```
eventmonitor -portal -migrate -source cx380_123_223 -dest
cx380_123_227
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-portal -destroy

Deletes the host agent from the portal system.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -portal -destroy` command deletes the specified host agent from the specified portal system.

SYNTAX

```
eventmonitor -portal -destroy -system portal system name
-server hostname [-o]
```

OPTIONS

`-system` *portal system name*

Specifies the portal system whose host is deleted.

`-server hostname`
 Deletes the specified host agent.

`-o`
 Executes the command without prompting for confirmation.

EXAMPLE

```
eventmonitor -portal -destroy -system cx380_123_227 -server
10.32.123.205
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-template -list

Lists the templates.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-list` function lists all the templates available in the template database of the target system.

SYNTAX

```
eventmonitor -template -list
```

OUTPUT

The templates in the database are listed; none if the templates are not available.

-template -import

Imports the template files.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-import` function imports the template file from the file system to the template database.

SYNTAX

```
eventmonitor -template -import -templatefilename filename
```

OPTIONS

`-templatefilename filename`
 Specifies the template to be imported to the template database of the target system.

EXAMPLE

```
eventmonitor -template -import -templatefilename
D:\Templates\Template_Test.tpl
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-template -export

Exports the templates.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-export` function exports the template from the template database to the file system.

SYNTAX

```
eventmonitor -template -export -templatename templatename
-filepath localfilepath
```

OPTIONS

```
-templatename template name
    Specifies the template.
-filepath localfilepath
    Specifies the local file path.
```

EXAMPLE

```
eventmonitor -template -export -templatename Template_Test
-filepath D:\Templates\
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-template -destroy

Deletes the templates.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-destroy` function deletes the specified template from the template database on the target system.

SYNTAX

```
eventmonitor -template -destroy -templatename templatename [-o]
```

OPTIONS

```
-templatename templatename
    Deletes the specified template.
-o
    Executes the command without prompting for confirmation.
```

EXAMPLE

```
eventmonitor -template -destroy -templatename Template_Test
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-template -findconflict

Finds and resolves template conflicts.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-findconflict` function lists the conflicting templates. The `-resolve` switch prompts you to resolve them.

SYNTAX

```
eventmonitor -template -findconflict [-resolve]
```

OPTIONS

`-resolve`

Resolves the template conflict.

EXAMPLE

```
eventmonitor -template -findconflict -resolve
```

OUTPUT

The sample output is as follows:

```
Event Template <template name> from host <host name> is
in conflict with the template in the database.
How do you want to resolve the conflict?
1. Update the template on the remote host.
2. Rename the template on the remote host.
3. Delete the template on the remote host.
4. Ignore this conflict.
Would you like to apply this option to all conflicting
template(s)? (y/n)
```

-template -swap

Swaps the two templates.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli eventmonitor -template` command with the `-swap` function swaps the two specified templates. The system that used template A uses template B after the swap.

SYNTAX

```
eventmonitor -template -swap -templateName template name A
template name B [-keepResponses]
```

OPTIONS

`-templateName` *template name A template name B*

Swaps the specified template names.

EXAMPLE

```
EventMonitor -template -swap -templateName Template_Test_1
Template_Test_2 -keepResponses
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -addsystem

Adds the system to be monitored.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-addsystem` function adds the specified system to the monitored system list of the specified host agent.

SYNTAX

```
eventmonitor -monitor -addsystem -server hostname -system
system name
```

OPTIONS

`-server hostname`
Specifies the host agent.

`-system system name`
Specifies the system.

EXAMPLE

```
eventmonitor -monitor -addsystem -server 10.32.123.205
-system cx380_123_223
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -removesystem

Removes the system from the monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-removesystem` function removes the specified system from the monitored system list of the specified host agent.

SYNTAX

```
eventmonitor -monitor -removesystem -server hostname -system
system name
```

OPTIONS

`-server hostname`
Specifies the host agent.

`-system system name`

Specifies the system.

EXAMPLE

```
eventmonitor -monitor -removesystem -server 10.32.123.205
-system cx380_123_223
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -listmapping

Lists the storage template mappings.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-listmapping` function lists all the storage template mappings of the specified centralized or distributed monitor.

SYNTAX

```
eventmonitor -monitor -listmapping -server hostname|-system
storage system name
```

OPTIONS

`-server hostname`

Lists the mapping information of the centralized monitor.

`-system system name`

Lists the mapping information of the distributed monitor.

EXAMPLE

```
eventmonitor -monitor -listmapping -server 10.32.123.205
```

```
eventmonitor -monitor -listmapping -system cx380_123_223
```

OUTPUT

The mapping information of the specified monitor is listed. None will be listed if the mapping is not available.

-monitor -applytemplate

Applies the template to the monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-applytemplate` function applies the template to the specific event monitor.

SYNTAX

```
eventmonitor -monitor -applytemplate [-server hostname]
[-system system name] -templatename templatename
```

OPTIONS

`-server hostname`

Applies the template to the specified centralized monitor as a global template.

`-system system name`

Applies the template to the specified distributed monitor. If you specify both the server and the system, it applies the template to the specific system monitored by the centralized monitor.

`-templatename templatename`

Specifies the template name.

EXAMPLE

```
eventmonitor -monitor -applytemplate -server 10.32.123.205
-system cx380_123_223 -templatename Template_Test
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -stoptemplate

Stops applying the template to the monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-stoptemplate` function stops applying the template to the specific event monitor.

SYNTAX

```
eventmonitor -monitor -stoptemplate [-server hostname] [-system
system name] -templatename templatename
```

OPTIONS

`-server hostname`

Stops applying the template to the specified centralized monitor.

`-system system name`

Stops applying the template to the specified distributed monitor. If you specify both the server and the system, it stops applying the template to the specific system and the centralized monitor.

`-templatename templatename`

Specifies the template name.

EXAMPLE

```
eventmonitor -monitor -stoptemplate -server 10.32.123.205 -system
cx380_123_223 -templatename Template_Test
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -getlog

Displays the event monitor log.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-getlog` function displays the event monitor log. If the range of entries is not specified, the entire log is displayed with the oldest entry first.

SYNTAX

```
eventmonitor -monitor -getlog -server hostname [-h] [+N] | [-N]
```

OPTIONS

`-server hostname`

Specifies the host agent.

`-h`

Displays the getlog header.

`+N`

Displays the oldest *n* entries in the log, with the oldest entry first.

`-N`

Displays the newest *n* entries in the log, with the newest entry first.

OUTPUT

```
Audit Logging Service
01/27/2009 14:41:46 N/A (4612)NaviCLI.exe application was
started by the user, following was the application detail
Process Id : 2540
Path :C:\emc\Navisphere\6.28.20.1.5.1\msgbin\NaviCLI.exe
User : user1
Domain : cx420_47_52.

Audit Logging Service
01/27/2009 14:41:46 N/A (2004)Test Event
NaviEventMonitorApp
01/27/2009 14:41:48 N/A (4613)NaviCLI.exe application was
closed by the user, following was the application detail
Process Id : 2540
Path :C:\emc\Navisphere\6.28.20.1.5.1\msgbin\NaviCLI.exe
User : user1
Domain : cx420_47_52.
```

-monitor -enableresponse

Enables the response from the specified monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-enableresponse` function enables the response from the specified event monitor.

SYNTAX

```
eventmonitor -monitor -enableresponse -server hostname |
-system system name
```

OPTIONS

```
-server hostname
    Enables the response from the centralized monitor.
-system system name
    Enables the response from the distributed monitor.
```

EXAMPLE

```
eventmonitor -monitor -enableresponse -system cx380_123_223
```

```
eventmonitor -monitor -enableresponse -server 10.32.123.205
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -disableresponse

Disables the response from the specified monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-disableresponse` function disables the response from the specified event monitor.

SYNTAX

```
eventmonitor -monitor -disableresponse -server hostname |
-system system name -time duration min [-o]
```

OPTIONS

```
-server hostname
    Disables the response from the centralized monitor.
-system system name
    Disables the response from the distributed monitor.
-time
    Specifies the time in minutes.
-o
    Executes the command without prompting for confirmation.
```

EXAMPLE

```
eventmonitor -monitor -disableresponse -system cx380_123_223
-time 30
```

```
eventmonitor -monitor -disableresponse -server 10.32.123.205
-time 30
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -getlogsize

Gets the log size of the monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-getlogsize` function gets the log size of the event monitor.

SYNTAX

```
eventmonitor -monitor -getlogsize -server hostname | -system
system name
```

OPTIONS

- `-server hostname`
Gets the log size of the centralized monitor.
- `-system system name`
Gets the log size of the distributed monitor.

EXAMPLE

```
eventmonitor -monitor -getlogsize -server 10.32.123.205
```

```
eventmonitor -monitor -getlogsize -system cx380_123_223
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -setlogsize

Sets the log size of the monitor.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-setlogsize` function changes the log size of Event Monitor.

SYNTAX

```
eventmonitor -monitor -setlogsize -server hostname | -system
system name -size logsize [-o]
```

OPTIONS

- `-server hostname`
Sets the log size of the centralized monitor.
- `-system system name`
Sets the log size of the distributed monitor.
- `-size logsize`
Specifies the log size in bytes.
- `-o`
Executes the command without prompting for confirmation.

EXAMPLE

```
eventmonitor -monitor -setlogsize -server 10.32.123.205
-size 20000000
```

```
eventmonitor -monitor -setlogsize -system cx380_123_223
-size 20000000
```

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -inserttestevent

Inserts an event in the log.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-inserttestevent` function inserts an event into the Event Monitor event log.

SYNTAX

```
eventmonitor -monitor -inserttestevent -server hostname
```

OPTIONS

`-server hostname`
Specifies the host agent.

OUTPUT

None if the command succeeds; status or error information if it fails.

-monitor -responsetest

Creates an artificial event.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-responsetest` function creates an artificial event to test the Event Monitor response.

SYNTAX

```
eventmonitor -monitor -responsetest
-server hostname
-email -smtpmailserver smtpmailserver -destaddress destaddress
[-subject subject] [-cc cc] [-message message|-file filename]
[-sender sender_email_addr]

-emailpage -smtpmailserver smtpmailserver -destaddress
destaddress [-subject subject] [-cc cc]
[-message message|-file filename] [-sender sender_email_addr]

-modempage -destnumber destnumber -msgnumber msgnumber
```

```
[-comport comport] [-dialcommand dialcommand]
[-initcommand initcommand] [-messagedelay messagedelay]

-snmp -desthost desthost [-device device] [-community
community]
```

OPTIONS

```
-server hostname
    Specifies the host agent.
-smtpmailserver smtpmailserver
    Specifies the SMTP mail server.
-destaddress destaddress
    Specifies the destination email address.
[-subject subject]
    Specifies the subject of the email.
[-cc cc]
    The email address to which to send a cc (copy).
[-message message|-file filename]
    The message describes the event and -file specifies a file with descriptive text.
[-sender sender_email_address]
    Specifies the sender address. If omitted, the CLI inserts the agent hostname as the
    sender_email_address.
-destnumber destnumber
    Specifies the pager phone number (with area code).
-msgnumber msgnumber
    Specifies the number that will appear on the pager display (with area code).
[-comport comport]
    Specifies the PC COM port number. If omitted, the CLI assumes COM 1.
[-dialcommand dialcommand]
    Specifies the modem dial command to be used by the agent.
[-initcommand initcommand]
    Specifies the modem dial command to be used by the agent. It allows for some
    custom modem configuration before the data is sent to it.
[-messagedelay messagedelay]
    Specifies the delay between sending the destination phone number to the modem
    and sending the message to the modem. This allows enough time for the receiving
    modem to answer the call.
-desthost desthost
    Specifies an SNMP management hostname.
-device device
    Specifies the device represented by the SNMP management host.
-community community
    Specifies a community on the SNMP management host.
```

OUTPUT

Provides an output depending on the switches used; provides status or error information if it fails.

-monitor -reloadconfig

Reloads the template file.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -monitor` command with the `-reloadconfig` function reloads the template file.

SYNTAX

```
eventmonitor -monitor -reloadconfig -server hostname
```

OPTIONS

`-server hostname`
Specifies the host agent.

OUTPUT

None if the command succeeds; status or error information if it fails.

Virtual server commands

This section explains the virtual server commands to configure virtual center/ESX credentials.

Virtual server overview

The virtual server integration feature simplifies the process for assigning a system to the virtual servers and their virtual machines.

The virtual server commands identify and display the ESX server or Virtual Center (VC) server, and display the mapping between the LUN and the virtual machine. These commands let you configure VC/ESX credentials.

You must configure an ESX server that needs virtualization integration information with a system in the domain. To execute host agent commands, the host does not need be attached to the system, but network connectivity is required.

[naviseccli virtual server commands on page 285](#) lists the virtual server commands.

Table 16 naviseccli virtual server commands

Command	Description
<code>server -volmap</code>	Lists the volume map information.
<code>server -update</code>	Updates the server information.
<code>server -getagent</code>	Displays the agent information.
<code>server -remoteconfig</code>	Displays the remote configuration information.
<code>server -register</code>	Displays the agent information on the host.
<code>hypervisor -addvirtualcenter</code>	Adds virtual center credentials.
<code>hypervisor -modifyvirtualcenter</code>	Modifies virtual center credentials.

Table 16 naviseccli virtual server commands (continued)

Command	Description
hypervisor -removevirtualcenter	Removes virtual center credentials.
hypervisor -addESX	Adds ESX credentials.
hypervisor -modifyESX	Modifies ESX server properties.
hypervisor -removeESX	Removes ESX credentials.
hypervisor -listESX	Lists ESX servers.

server -volmap

Lists the volume map information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -volmap` command lists the volume map information. If the host is an ESX server (ESX_1), then it displays the ESX server information (LUN to VM mappings) of all the systems attached to the host. If the specified host is a physical host, it displays `lunmapinfo` information for all systems attached to the host. Refer to the *EMC Navisphere CLI Reference* (P/N 300-003-628 Rev. A10) for information on the `lunmapinfo` command.

SYNTAX

```
naviseccli -h array_a server -volmap -host ipaddress [-local] [-vm]
```

OPTIONS

`-host ipaddress`

Specifies the IP address of the attached host whose information is displayed.

`-local`

Displays only the attached host information.

`-vm`

Displays all ESX_1 mount points on systems A and B.

EXAMPLE

```
naviseccli server -volmap -host ipaddress
```

OUTPUT

Sample output (if the host is a physical host):

```
Logical Drives:          N\A
Physical Device:        \\.\PhysicalDrive13
LOGICAL UNIT WWN:      60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:11
LOGICAL UNIT NUMBER:   3
Current Owner:         N\A
Default Owner:         N\A
Trespassed:            N\A
LUN Capacity(Gigabytes): .2
LUN Capacity(Blocks):  4194304
```

```

Logical Drives:          P:\
Physical Device:         \\.\PhysicalDrive14
LOGICAL UNIT WWN:
60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:12
LOGICAL UNIT NUMBER:    4
Current Owner:          SP A
Default Owner:          SP B
Trespassed:             Yes
LUN Capacity(Gigabytes): 1.2
LUN Capacity(Blocks):   2097152
naviseccli server -volmap -host ipaddress -local

Logical Drives:          F:\
Physical Device:         \\.\PhysicalDrive18
LOGICAL UNIT WWN:
60:06:01:60:36:F7:22:00:2A:48:A4:C1:57:0E:DE:11
SNAPSHOT admhost_snapshot1_0 ----- snapshot name
will replace LUN
Default Owner:          SP A
Default Owner:          SP A
Trespassed:             NO
LUN Capacity(Gigabytes): 0
LUN Capacity(Blocks):   204800

```

Sample output (if the host is an ESX server):

```

ESX Server Info:
Name: nlpc12197.us.dg.com
IP address: 10.14.12.197

Device Id: vmhba1:0:1
Device Name: vmhba1:0:1

Datastore:N/A
LOGICAL UNIT WWN:
60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:11
LOGICAL UNIT NUMBER:    4
Current Owner:          SP A
Default Owner:          SP B
LUN Capacity(Gigabytes): .2
LUN Capacity(Blocks):   4194304
naviseccli server -volmap -host ipaddress -vm

```

Sample output (if the host is an ESX server):

```

ESX Server Info:
Name: nlpc12197.us.dg.com
IP address: 10.14.12.197

Virtual Machine Info:
VM Name: Hard Disk 1
Guest Host Name:
Guest IP Address:
Guest OS:

Name: Hard Disk 1
Type: Mapped Raw LUN
File Size(MB):100
File Path: \[DataStore1]\VM1\hd1.vmdk
LUN Name: LUN 5
LUN ID: 5
Name: My virtual machine
Type: VMX Config File
File Size(MB): N/A
File Path: \[DataStore1]\VM1\hd1.vmdk

```

```
LUN Name: N/A
LUN ID: N/A
```

server -update

Updates the server information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -update` command polls and updates the physical and ESX servers attached to the system. It also displays the update status.

SYNTAX

```
naviseccli -h sp server-update [-host ipaddress
[-rescandevices[-o]]]
[-all [-rescandevices[-o]]] [-status]
```

OPTIONS

`-host ipaddress`

Specifies the IP address of the host (physical host or ESX server). It polls and updates the host.

`-all`

Polls and updates all the hosts (physical and ESX server) attached to the system.

`-status`

Displays the status of all hosts polled in the last poll cycle.

`-rescandevices`

Rescans and then updates the servers. It displays the confirmation before proceeding.

`-o`

Overrides the confirmation.

EXAMPLE

```
naviseccli server -update -status
```

OUTPUT

```
Update Status:
Poll Servers started at 06-NOV-2008 16:19:53.
Poll Servers completed successfully at 06-NOV-2008 16:22:53
Processed 3 out of 3 servers
```

server -getagent

Gets the agent information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -getagent` command lists the agent information on the host.

SYNTAX

```
naviseccli -h sp server -getagent -host ipaddress
```


OPTIONS

`-host ipaddress`

Specifies the IP address of an attached host. It is applicable for any host running the agent.

OUTPUT

Agent Rev: 6.29.75 (0.9)

server -remoteconfig

Gets the remote configuration information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -remoteconfig` command displays and modifies the host agent configuration information. The command scans all the devices in the system on a remote host.

SYNTAX

```
naviseccli -h sp server -remoteconfig -host ipaddress
```

OPTIONS

`-host ipaddress`

Specifies the IP address of an attached host. It is applicable for any host running the agent.

`-description`

Gets/sets the contents of description field.

`-contact`

Gets/sets the contents of the contact field.

`-dev`

Displays the device name, system, comments, and the connection type (SCSI, TTY, and LAN).

`-users`

Displays information about the users.

`-interval`

Gets/sets the contents of the polling interval field. The valid range is 1 to 120.

`-baudrate`

Gets/sets the baud rate.

`-logsize`

Gets/sets the log size to be transferred.

`-userexplicitdevnames`

Gets/sets user explicit device names.

`-write`

Writes an image of the configuration file as it exists on the host.

`-managedev`

Adds the SCSI devices in the managed device list. Each device includes a device name, system name, and an optional description.

`-open`

Keeps the device handle open between requests. This improves or degrades the agent performance depending on the operating system.

`-managetty`

Adds the serial devices in the managed device list. Each device includes a device name, system name, and an optional description.

`-managelan`

Adds the LAN devices in the managed device list. Every device includes a device name, system name, and an optional description.

`-adduser`
Sets the users field. The format of the usernames for SP A/SP B is `system@ipaddress of SPA/SPB` ; separate these usernames with a space.

`-unmanageddev`
Removes the device from the managed device list.

`-manageauto`
Executes auto configuration (manages all SCSI devices).

`-rmuser`
Removes the specified users from the authorized user list.

`-f`
Reads an image of the configuration file as it exists on the host.

`-reloadconfig`
Reloads the agent configuration file without stopping and starting the agent. Once polling of the server has completed, the changes will be reflected.

`-scan`
Scans the devices in the system on a remote host.

OUTPUT

Provides output depending on the options used.

server -register

Calls the Unisphere Agent running on the server to push/registration to systems.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -register` command calls the Unisphere Agent running on the server to push/registration to systems.

SYNTAX

```
naviseccli -h sp server -register -host ipaddress
```

OPTIONS

`-host ipaddress`

Specifies the IP address of an attached host. It is applicable for any host running the agent.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -addvirtualcenter

Adds virtual center credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -addvirtualcenter` command adds the virtual center credentials. The virtual center credentials are stored in the system.

SYNTAX

```
naviseccli -h sp server -hypervisor  
-addvirtualcenter ipaddress
```

```
-username user [-password password]
[-description description]
[-cert pathToCertificatefile]
[-o]
```

OPTIONS

ipaddress

Specifies the IP address of the virtual center server.

-username *user*

Specifies the username of the virtual center.

-password *password*

Specifies the password of the virtual center. If you do not type the password, the system prompts you to type it.

-description *description*

Maximum variable length is 128 characters and the minimum variable length is 0.

-cert *pathToCertificatefile*

Specifies the full pathname of trusted certificate file to be uploaded to the certificate store.

-o

Sends a prompt if you specify the incorrect username/password.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -modifyvirtualcenter

Modifies virtual center credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -modifyvirtualcenter` command modifies an existing virtual center server's credentials.

SYNTAX

```
naviseccli -h sp server -hypervisor -modifyvirtualcenter name |
ipaddress [-username user -password password] [-o]
[-description description]
```

OPTIONS

name | *ipaddress*

Specifies the name or the IP address of the virtual center server.

-username *user*

Specifies the new username of the virtual center.

-password *password*

Specifies the password of the virtual center.

-description

Maximum variable length is 128 characters and the minimum variable length is 0.

-o

Sends a prompt if you specify the incorrect username/password.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -removevirtualcenter

Removes virtual center credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -removevirtualcenter` command removes an existing virtual center server's credentials.

SYNTAX

```
naviseccli -h sp server -hypervisor -removevirtualcenter name |
ipaddress [-o]
```

OPTIONS

`name | ipaddress`

Specifies the name or the IP address of the virtual center server.

`-o`

Executes the command without prompting for confirmation.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -addesx

Adds ESX credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -addesx` command adds ESX credentials. You must specify the ESX credentials manually if the ESX server is not managed by the virtual center server.

SYNTAX

```
naviseccli -h sp server -hypervisor -addesx name|ipaddress
-username user [-password password] [-cert
pathToCertificatefile] [-o]
```

OPTIONS

`-username user`

Specifies the username of the ESX server.

`-password password`

Specifies the password of the ESX server.

`-cert pathToCertificatefile`

Specifies the full pathname of the trusted certificate file to be uploaded to the certificate store.

`-o`

Sends a prompt if you specify the incorrect username/password.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -modifyesx

Modifies ESX server properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -modifyesx` command modifies an existing virtual center server's credentials.

SYNTAX

```
naviseccli -h sp server -hypervisor -modifyesx name | ipaddress
-username user -password password [-o]
```

OPTIONS

`name | ipaddress`

Specifies the name or the IP address of the ESX server to be modified.

`-username user`

Specifies the new username of the ESX server.

`-password password`

Specifies the new password of the ESX server.

Note

The options `-username` and `-password` have to be specified to change either one of them.

`-o`

Sends a prompt if you specify the incorrect username/password.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -removeesx

Removes ESX credentials.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -removeesx` command removes existing ESX server credentials.

SYNTAX

```
naviseccli -h sp server -hypervisor -removeesx name | ipaddress
[-o]
```

OPTIONS

`name | ipaddress`

Specifies the name or the IP address of the ESX server to be removed.

`-o`

Executes the command without prompting for confirmation.

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -listesx

Lists ESX servers.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -listesx` command lists all servers with IP address and description.

SYNTAX

```
naviseccli -h sp server -hypervisor -listesx
```

OUTPUT

```
ESX Server(s):
Name: nlpc12241.us.dg.com
IP Address: 10.14.12.241
Virtual center Managed ESX server(s):
Virtual Center IP Address: 10.14.12.82
Virtual Center Description:
Virtual Center IP Address: 10.14.12.248
Virtual Center Description:
```

hypervisor -bypassCert -set

Sets bypass certificate verification.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -bypassCert -set` command sets the bypass certificate verification for virtual servers.

SYNTAX

```
naviseccli -h sp server -hypervisor -bypassCert -set [0|1]
```

OPTIONS

`-set`

Sets the bypass certificate verification for virtual servers. Valid values are 0 (not bypass) and 1 (bypass).

OUTPUT

None if the command succeeds; status or error information if it fails.

hypervisor -bypassCert -get

Gets bypass certificate verification.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli hypervisor -bypassCert -get` command gets the bypass certificate verification for virtual servers.

SYNTAX

```
naviseccli -h sp server -hypervisor -bypassCert -get
```

OPTIONS`-get`

Gets the bypass certificate verification for virtual servers.

OUTPUT

None if the command succeeds; status or error information if it fails.

Host Agent commands

This section explains the host agent commands to configure virtual center/ESX credentials.

server -volmap

Lists the volume map information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -volmap` command lists the volume map information. If the host is an ESX server (ESX_1), then it displays the ESX server information (LUN to VM mappings) of all the systems attached to the host. If the specified host is a physical host, it displays lunmapinfo information for all systems attached to the host. Refer to the *EMC Navisphere CLI Reference* (P/N 300-003-628 Rev. A10) for information on the `lunmapinfo` command.

SYNTAX

```
naviseccli -h array_a server -volmap -host ipaddress [-local]
[-vm]
```

OPTIONS`-host ipaddress`

Specifies the IP address of the attached host whose information is displayed.

`-local`

Displays only the attached host information.

`-vm`

Displays all ESX_1 mount points on systems A and B.

EXAMPLE

```
naviseccli server -volmap -host ipaddress
```

OUTPUT

Sample output (if the host is a physical host):

```
Logical Drives:          N\A
Physical Device:         \\.\PhysicalDrive13
LOGICAL UNIT WWN:        60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:11
LOGICAL UNIT NUMBER:     3
Current Owner:           N\A
Default Owner:           N\A
Trespassed:              N\A
LUN Capacity(Gigabytes): .2
LUN Capacity(Blocks):    4194304
Logical Drives:          P:\
```

```

Physical Device:          \\.\PhysicalDrive14
LOGICAL UNIT WWN:
60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:12
LOGICAL UNIT NUMBER:    4
Current Owner:          SP A
Default Owner:          SP B
Trespassed:             Yes
LUN Capacity(Gigabytes): 1.2
LUN Capacity(Blocks):   2097152
naviseccli server -volmap -host ipaddress -local

Logical Drives:         F:\
Physical Device:        \\.\PhysicalDrive18
LOGICAL UNIT WWN:
60:06:01:60:36:F7:22:00:2A:48:A4:C1:57:0E:DE:11
SNAPSHOT admhost_snapshot1_0 ----- snapshot name
will replace LUN
Default Owner:          SP A
Default Owner:          SP A
Trespassed:             NO
LUN Capacity(Gigabytes): 0
LUN Capacity(Blocks):   204800

```

Sample output (if the host is an ESX server):

```

ESX Server Info:
Name: nlpc12197.us.dg.com
IP address: 10.14.12.197

Device Id: vmhba1:0:1
Device Name: vmhba1:0:1

Datastore:N/A
LOGICAL UNIT WWN:
60:06:01:60:93:31:19:00:01:DC:1D:0B:07:0D:DB:11
LOGICAL UNIT NUMBER:    4
Current Owner:          SP A
Default Owner:          SP B
LUN Capacity(Gigabytes): .2
LUN Capacity(Blocks):   4194304
naviseccli server -volmap -host ipaddress -vm

```

Sample output (if the host is an ESX server):

```

ESX Server Info:
Name: nlpc12197.us.dg.com
IP address: 10.14.12.197

Virtual Machine Info:
VM Name: Hard Disk 1
Guest Host Name:
Guest IP Address:
Guest OS:

Name: Hard Disk 1
Type: Mapped Raw LUN
File Size(MB):100
File Path: \[DataStore1]\VM1\hd1.vmdk
LUN Name: LUN 5
LUN ID: 5
Name: My virtual machine
Type: VMX Config File
File Size(MB): N/A
File Path: \[DataStore1]\VM1\hd1.vmdk
LUN Name: N/A
LUN ID: N/A

```


server -update

Updates the server information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -update` command polls and updates the physical and ESX servers attached to the system. It also displays the update status.

SYNTAX

```
naviseccli -h sp server-update [-host ipaddress
[-rescandevices[-o]]]
[-all [-rescandevices[-o]]] [-status]
```

OPTIONS

`-host ipaddress`

Specifies the IP address of the host (physical host or ESX server). It polls and updates the host.

`-all`

Polls and updates all the hosts (physical and ESX server) attached to the system.

`-status`

Displays the status of all hosts polled in the last poll cycle.

`-rescandevices`

Rescans and then updates the servers. It displays the confirmation before proceeding.

`-o`

Overrides the confirmation.

EXAMPLE

```
naviseccli server -update -status
```

OUTPUT

```
Update Status:
Poll Servers started at 06-NOV-2008 16:19:53.
Poll Servers completed successfully at 06-NOV-2008 16:22:53
Processed 3 out of 3 servers
```

server -getagent

Gets the agent information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -getagent` command lists the agent information on the host.

SYNTAX

```
naviseccli -h sp server -getagent -host ipaddress
```

OPTIONS

`-host ipaddress`

Specifies the IP address of an attached host. It is applicable for any host running the agent.

OUTPUT

Agent Rev: 6.29.75 (0.9)

server -remoteconfig

Gets the remote configuration information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -remoteconfig` command displays and modifies the host agent configuration information. The command scans all the devices in the system on a remote host.

SYNTAX

```
naviseccli -h sp server -remoteconfig -host ipaddress
```

OPTIONS

`-host ipaddress`

Specifies the IP address of an attached host. It is applicable for any host running the agent.

`-description`

Gets/sets the contents of description field.

`-contact`

Gets/sets the contents of the contact field.

`-dev`

Displays the device name, system, comments, and the connection type (SCSI, TTY, and LAN).

`-users`

Displays information about the users.

`-interval`

Gets/sets the contents of the polling interval field. The valid range is 1 to 120.

`-baudrate`

Gets/sets the baud rate.

`-logsize`

Gets/sets the log size to be transferred.

`-userexplicitdevnames`

Gets/sets user explicit device names.

`-write`

Writes an image of the configuration file as it exists on the host.

`-managedev`

Adds the SCSI devices in the managed device list. Each device includes a device name, system name, and an optional description.

`-open`

Keeps the device handle open between requests. This improves or degrades the agent performance depending on the operating system.

`-managetty`

Adds the serial devices in the managed device list. Each device includes a device name, system name, and an optional description.

`-managelan`

Adds the LAN devices in the managed device list. Every device includes a device name, system name, and an optional description.

`-adduser`

Sets the users field. The format of the usernames for SP A/SP B is `system@ipaddress of SPA/SPB ; separate these usernames with a space.`

`-unmanageddev`
Removes the device from the managed device list.

`-manageauto`
Executes auto configuration (manages all SCSI devices).

`-rmuser`
Removes the specified users from the authorized user list.

`-f`
Reads an image of the configuration file as it exists on the host.

`-reloadconfig`
Reloads the agent configuration file without stopping and starting the agent. Once polling of the server has completed, the changes will be reflected.

`-scan`
Scans the devices in the system on a remote host.

OUTPUT

Provides output depending on the options used.

server -register

Calls the Unisphere Agent running on the server to push/registration to systems.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli server -register` command calls the Unisphere Agent running on the server to push/registration to systems.

SYNTAX

```
naviseccli -h sp server -register -host ipaddress
```

OPTIONS

`-host ipaddress`
Specifies the IP address of an attached host. It is applicable for any host running the agent.

OUTPUT

None if the command succeeds; status or error information if it fails.

Unisphere Server Utility

This section provides the CLI commands for Unisphere Server Utility tools.

Using the server utility command line to generate a high-availability report

The high-availability option determines if the server is configured for high availability by verifying that the server has at least one connection path to each system SP, and that PowerPath or some other failover software, such as DMP, VMware native, PV Links, or HP native failover is running. The utility will not detect any other native failover software, such as native multipath (MPIO) for Linux, Solaris, or Windows Server 2008.

Generate a high-availability report to:

- check the status of a server.
- prepare for software installation on a system.

Checking the status of a server

To periodically check the high-availability status of a server, issue the appropriate command for your operating system.

For UNIX servers:

```
naviserverutilcli hav[-directory | -d outputdirectory
```

where *outputdirectory* specifies the directory to which you want to deposit the report on the server. This switch is required only if you are running the server utility from the CD. Otherwise the report is automatically saved to the Unisphere Server Utility installation directory. If the server utility is installed on your server and you specify a directory, the directory you specify will override the server utility's installation directory.

For Windows servers:

1. Open a command window and enter one of the following:

- `cd C:\Program Files\EMC\Unisphere Server Utility`
- `cd C:\Program Files (x86)\EMC\Unisphere Server Utility`

2. Enter:

```
naviserverutilcli hav[-directory | -d outputdirectory
```

where *outputdirectory* specifies the directory to which you want to deposit the report on the server. This switch is required only if you are running the server utility from the CD. Otherwise the report is automatically saved to the Unisphere Server Utility installation directory. If the server utility is installed on your server and you specify a directory, the directory you specify will override the server utility's installation directory.

3. For ESX servers, use the following additional switches to specify the ESX server or vCenter for which you want a report.

- `-esx ipaddress`
Specifies the IP address of the VMware ESX Server's environment for which you want to generate a report.
- `-virtualcenter ipaddress`
Specifies the IP address of the vCenter. You must also specify the `-esx ipaddress` for which you want to generate a report.
- `-vmuser username`
Specifies the username for the ESX Server or vCenter.
- `-vmpassword password`
Specifies the password for the ESX Server or vCenter.

Preparing for software installation on a system

If you are generating a high-availability report as a prerequisite for updating the system software, you need to upload this report to the system prior to installing the software on the system.

To generate and upload an HA report, issue the following command from the directory in which you installed the utility:

```
naviserverutilcli hav -upload -ip IPAddress [-h
IPAddress|network_name -user|-u username
-password|-p userpassword -scope|-s scopelevel -secfilepath
securityfilepath -port 443|2163[-directory|-d outputdirectory ]
-esx ipaddress -virtualcenter -virtualcenter ipaddress -vmuser
username -vmpassword password
-ip IPAddress
```

- Specifies the IP address of an SP on the system to which you will upload the report. You cannot use this switch in conjunction with the `-h` switch.
- `-h IPaddress | network_name`
Specifies the IP address or the network name of an SP on the system to which you will upload the report. You cannot use this switch in conjunction with the `-ip` switch.
- `-user|-u username`
Specifies the login name for the system. You must use this switch in conjunction with the `-password|-p` and `-scope|-s` switches. You cannot use this switch in conjunction with the `-secfilepath` switch.
- `-password|-p userpassword`
Specifies the password required to access the system. You must use this switch in conjunction with the `-user|-u` and `-scope|-s` switches. You cannot use this switch in conjunction with the `-secfilepath` switch.
- `-scope|-s scopelevel`
Specifies the level of access the username and password has to the system: 0 = global; 1 = local; 2 = LDAP. You must use this switch in conjunction with the `-user|-u` and `-password|-p` switches. You cannot use this switch in conjunction with the `-secfilepath` switch.
- `-secfilepath securityfilepath`
Searches the security file for the specified IP address. If the utility does not find the specified IP address, then the utility will use the default entry in the security file. A security file is required to run this switch. For information on how to create a security file, refer to the EMC Navisphere Command Line (CLI) Reference guide. You cannot use this switch in conjunction with the `-user|-u`, `-password|-p`, and `-scope|-s` switches.
- `-port 443|2163`
Specifies the Unisphere Management Server port number. The default port is 443.
- `-directory|-d outputdirectory`
Specifies the directory to which you want to deposit the report on the server. This switch is required only if you are running the server utility from the CD. Otherwise the report is automatically saved to the Unisphere Server Utility installation directory. If the server utility is installed on your server and you specify a directory, the directory you specify will override the server utility's installation directory.
- `-esx ipaddress`
Specifies the IP address of the VMware ESX Server's environment for which you want to generate a report.
- `-virtualcenter ipaddress`
Specifies the IP address of the Virtual Center. You must also specify the `-esx ipaddress` for which you want to generate a report.
- `-vmuser username`
Specifies the username for the ESX Server or Virtual Center.
- `-vmpassword password`
Specifies the password for the ESX Server or Virtual Center.

After uploading the HA report

In addition to running the server utility to validate server high availability, we strongly recommend that you perform the following manual checks:

Verify support of software and hardware

Make sure that you have verified that all software and hardware are supported according to the EMC E-Lab™ Interoperability Navigator and support matrices. This

tool does not verify that you are running supported switch firmware, nor that you are using redundant switches (recommended).

Upgrading system software

If you are about to perform a system online software upgrade or you added an HBA, a LUN, or changed connectivity (added a path or changed zoning):

- Make sure that you have validated all software and hardware that are supported and interoperable according to EMC E-Lab Interoperability Navigator and support matrices.
- Verify that all LUNs (virtual disks) that your servers will access during the upgrade are under path management software control. For servers running PowerPath and DMP, see the Failover Software section on the Details tab of the report and locate these LUNs.

Unisphere Initialization Utility

This section provides the CLI commands for Unisphere Initialization Utility tools.

Using the command line initialization utility

For Fibre Channel systems, use the Unisphere Initialization Utility to discover systems and set network parameters (IP address, subnet mask, and default gateway). In addition, for iSCSI systems with iSCSI data ports attached to Windows server, use the utility to set network parameters for these ports.

You can install the utility on a server or other host such as an off-array management station. The only requirement is that the host be connected to the same network subnet as the system management ports.

After you start the command line version of the initialization utility, use the `naviinittoolcli` command with any of the switches below.

```
eula [-language language] discover [-all] [-xml]
configure -serial serialnumber [-file filename]
[-ipa spAIpaddress |-sphosta spAhostname] [-ipb spBIpaddress |
-sphostb spBhostname] [-mask subnetmask] [-gateway
gatewayaddress]
[-user username] [-password password] [-storagename
storageName]
[-ipv6mode |automatic |disable | [manual [-globalprefix prefix
[-ipv6gateway ipv6gateway]]] [-help]
```

where:

`eula`

Displays the EMC end-user license agreement (EULA) in the specified language.

`-language language`

The default language is English. Currently, the only valid language is English.

`discover`

Discovers and displays a list of partially initialized systems.

`-all`

Discovers and displays a list of initialized and partially initialized systems.

`-xml`

Displays the output in xml format.

`configure`

Specifies the network parameters for the specified system.

`-serial serialnumber`

Specifies the serial number of the system you want to initialize. If a system with the requested serial number is not found, the command fails.

- file *filename*
Specifies the name of an optional input file that stores all the network parameters. Both the SP management ports need to be on the same subnet for the inittool command to function appropriately. If you want the ports on different subnets, you can change subnets after the inittool has run.
- ipa *spAipaddress* | -sphosta *spAhostname*
Specifies the IP address of the management port on SP A, or the server name for SP A for the specified system.
- ipb *spBipaddress* | -sphostb *spBhostname*
Specifies the IP address of the management ports on SP B, or the server name for SP B for the specified system.
- mask *subnetmask*
Specifies the subnet mask for the specified system.
- gateway *gateway*
Specifies the gateway for the specified system.
- user *username*
Specifies the system login username.
- password *password*
Specifies the system login password.
- storagename *storagename*
Specifies the storage name for the specified system.
- help
Displays the help screen and does not start the initialization process.
- ipv6mode automatic | disable | [manual -globalprefix *prefix* | -
ipv6gateway *ipv6gateway*]
Specifies the configuration type for the IPv6 network. Refer to the IPv6 configuration type table for details.

Table 17 IPv6 configuration types

IPv6mode	Description
automatic	Automatically sets the IPv6 global prefix and IPv6 gateway for the interface, based on Router Advertisements.
disable	Disables the IPv6 functionality on this interface.
manual -globalprefix <i>prefix</i> [-ipv6gateway <i>gateway</i>]	Initializes the interface with the user specified prefix and (optional) gateway. It ignores any Router Advertisements.

CHAPTER 3

Virtual Pools Commands

This chapter provides a brief overview of the Virtual Pools feature and describes the storage pool and thin commands that you can use if this feature is enabled on your system.

Major topics include:

- [Virtual Pools overview](#)..... 306
- [storagepool -create](#)..... 307
- [storagepool -list](#)..... 310
- [storagepool -destroy](#)..... 316
- [storagepool -modify](#)..... 316
- [storagepool -expand](#)..... 318
- [storagepool -cancelexpand](#)..... 319
- [storagepool -feature -info](#)..... 319
- [storagepool -resumeharvesting](#)..... 321
- [lun -create](#)..... 321
- [lun -modify](#)..... 323
- [lun -expand](#)..... 325
- [lun -destroy](#)..... 326
- [lun -list](#)..... 326
- [deduplication -feature -pause](#)..... 330
- [deduplication -feature -resume](#)..... 330
- [deduplication -feature -info](#)..... 330

Virtual Pools overview

The Virtual Pools feature allows you to allocate storage using thick and thin LUNs within pools and requires that the thin provisioning enabler be installed on the system. Thick LUNs and thin LUNs can reside within the same pool and they share the pool's storage capacity. Both thick and thin LUNs can be easily provisioned, expanded, and compressed by using the Unisphere interface.

A thin LUN lets you assign more storage capacity to a host than is physically available. Storage is assigned to the server in a capacity-on-demand method from a shared pool. A thin LUN competes with other LUNs in the pool for the available pool storage. The storage system software monitors and adds storage capacity, as required, to each pool, not each LUN. This simplifies the creation and allocation of storage capacity.

Note

The actual consumed capacity and rate of consumption for the thin LUN can vary depending on the attached host file system or the application using the LUN. This is a normal condition typical of most thin provisioning services.

A thick LUN reserves all of its assigned capacity at the time it is created.

Current Virtual Pool restrictions

- Thin LUNs cannot be used in the reserved LUN pool, the Write Intent Log, or as Clone Private LUNs (CPLs).
- Thick LUNs cannot be used in the Write Intent Log or as CPLs. They can be included in the reserved LUN pool.
- Thin and thick LUNs cannot be a component in a metaLUN.
- Virtual Pools support RAID 5, RAID 6 and RAID 1/0 types. RAID 5 is the default RAID type for the Extreme Performance and Performance Tiers, and RAID 6 is the default RAID type for the Capacity Tier.
- Pools support Flash (Extreme Performance), SAS (Performance), and NL-SAS (Capacity) drives.

Pools

A pool contains a set of disks that can include any of the following drive types: Flash, SAS, and NL-SAS. The maximum number of disks for a pool depends on system type. For more efficient performance, we recommend that all disks in the pool have the same capacity and RAID type. RAID 5 is the default RAID type for the Extreme Performance and Performance Tiers, and RAID 6 is the default RAID type for the Capacity Tier. A storage pool shares its user capacity with all the LUNs in the pool. You can expand the user capacity of a pool by adding disks to it.

Note

Vault drives (the first four drives in a system) cannot be part of a pool. The Unisphere dialogs and wizards will not let you select these drives.

For the latest pool LUN limit information, refer to the product's *Simple Support Matrix* on [EMC Online Support](#).

Table 18 navisecli thin provisioning commands

Command	Description
Storage pool commands	
<code>storagepool -create</code>	Creates a thin pool.
<code>storagepool -list</code>	Lists the thin pools properties.
<code>storagepool -destroy</code>	Deletes a thin pool.
<code>storagepool -modify</code>	Modifies thin pool properties.
<code>storagepool -expand</code>	Expands the thin pool on the system.
<code>storagepool -cancelexpand</code>	Cancel the expansion, if the expansion has failed.
<code>storagepool -feature -info</code>	Lists out the parameters and settings of the mapped driver and certain configuration information.
<code>storagepool - resumeharvesting</code>	Resumes harvesting on an existing storage pool
LUN commands	
<code>lun -create</code>	Creates a pool LUN or a thin LUN depending if the <code>-type</code> switch is specified.
<code>lun -destroy</code>	Deletes a LUN.
<code>lun -list</code>	Lists the LUNs on a system and their properties.
<code>lun -modify</code>	Modifies the LUN properties.
<code>lun -expand</code>	Expands a pool LUN.
Deduplication commands	
<code>deduplication -feature -pause</code>	Pauses deduplication.
<code>deduplication -feature -resume</code>	Resumes deduplication.
<code>deduplication -feature -info</code>	Displays deduplication status.

storagepool -create

Creates the thin pool.

PREREQUISITES

For `navisecli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `navisecli storagepool -create` command creates the thin pool. A thin pool contains physical disks on which you bind thin LUNs. You can create more than one thin pool on a system, each with a different set of disks. There are two types of storage pools: Thin pools and RAID groups. A thin pool contains thin LUNs. If you specify the RAID type when creating the storage pool, the specified RAID type is used for creating the storage pool. If you do not specify the RAID type, the storage pools are created based on best practices. This may result in a storage pool with multiple RAID types. If you want to create a storage pool with 3 tiers and you specify the RAID types for each tier, then you will need to use the following steps:

1. Create the storage pool with the disks for the first tier and wait until it initializes.
2. Expand the storage pool with the disks for the second tier and wait for the expand to complete.
3. Expand the storage pool with the disks for the third tier.

SYNTAX

```
storagepool -create -disks diskslist [-rtype
raidType[-rdrivecount drivecount]]
[-description description]
[-name name] [-prcntFullThreshold threshold]
[-autoTiering scheduled|manual (default)]
[-fastcache on|off] [snapPoolFullThresholdEnabled on|off]
[-skipRules]
[-snapPoolFullHWM poolFullHWM] [-snapPoolFullLWM poolFullLWM]
[-snapSpaceUsedThresholdEnabled on|off]
[-snapSpaceUsedHWM snapSpaceUsedHWM]
[-snapSpaceUsedLWM snapSpaceUsedLWM]
[-initialverify yes|no] [-dedupRate Low|Medium|High]
[-dedupTieringPolicy noMovement|autoTier|highestAvailable|
lowestAvailable]
[-dedupInitialTier optimizePool|highestAvailable|
lowestAvailable]
```

OPTIONS

-disks

Specifies the list of disks in *Bus X Enclosure Y Disk Z* format. If you do not specify Bus in the disk format, Bus 0 is assumed.

-rtype

Specifies the type of RAID group of the thin pool. The valid RAID types are r_5, r_6, and r_10.

-rdrivecount

Specifies the RAID drive count used to create internal RAID groups.

-description

Sets the description for the thin pool to be created. The character length ranges from 0 to 255.

-name

Sets the unique name for the thin pool. The character length ranges from 1 to 64.

-prcntFullThreshold

Specifies the percentage of the thin pool to be used before the system generates the alerts. The percentage ranges from 1 to 84. The default value is 70%.

-skipRules

Skips the best practices check while creating the thin pools. If you do not specify this option, the thin pool creation may fail due to any error checked in the thin rules.

-autoTiering *scheduled|manual*

Sets the auto-tiering state to *scheduled* or *manual* for the storage pool. The default auto-tiering state is *manual*.

-fastcache *on|off*

Enables or disables the FAST Cache for the entire storage pool. The default value is enabled.

-snapPoolFullThresholdEnabled

Indicates that the system should check the pool full high water mark for auto delete. Valid values are On or Off.

-snapPoolFullHWM

Specifies the pool full high watermark that triggers the auto delete process of the storage pool to be created.

- snapPoolFullLWM
Specifies the pool full low watermark that stops the auto delete process of the storage pool to be created.
- snapSpaceUsedThresholdEnabled
Indicates that the system should check the snapshot space used high water mark for auto delete. Valid values are On and Off.
- snapSpaceUsedHWM
Specifies the snapshot space used high watermark that triggers the auto delete process of the storage pool to be created.
- snapSpaceUsedLWM
Specifies the snapshot space used low watermark that stops the auto delete process of the storage pool to be created.
- initialverify
Allows the user to specify not to have initial background verify during pool creation and expansion in order to improve operation performance of pool creation and expansion.
- dedupRate
Specifies the deduplication rate on this storage pool. Valid values are Low, Medium, or High.
- dedupTieringPolicy
Specifies the common FAST VP tiering policy for all deduplicated LUNs in this storage pool. Valid values are noMovement, autoTier, highestAvailable, or lowestAvailable.
- dedupInitialTier
Specifies the common initial tier for all deduplicated LUNs in this storage pool. Valid values are optimizePool, highestAvailable, or lowestAvailable.

EXAMPLE #1

This command creates a storage pool with a mix of Flash drives (4 – 8) and SAS drives (9 – 13). In this case, the best practices are used to create the storage pool for each tier by using RAID 5.

```
naviseccli -h riker-spa storagepool -create -disks 0_0_4 0_0_5
0_0_6 0_0_7 0_0_8 0_0_9 0_0_10 0_0_11 0_0_12 0_0_13 -rtype r_5
```

EXAMPLE #2

```
naviseccli -h riker-spa storagepool -create -disks 0_1_0 0_1_1 0_1_2
0_1_3 0_1_4 0_1_5 0_1_6 0_1_7 0_1_8 -rtype r_5 -rdrivecount 9
```

OUTPUT

None if the command succeeds; status or error information if initiating the command fails.

EXAMPLE #3

```
naviseccli -h riker-spa storagepool -create -disks 0_0 0_1 0_2 -rtype
r_5 -dedupRate high -dedupTieringPolicy autoTier -dedupInitialTier
highestAvailable
```

OUTPUT

None if the command succeeds; status or error information if initiating the command fails

If the user's choice of disks for the storage pool or RAID Group breaks the hot spare policy and `-skipRules` option is not used, the system will reject the request with the following message:

```
The current selection of disks will break the hot spare policy.
```

storagepool -list

Lists the thin pool properties and RAID groups.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -list` command lists the thin pool properties and RAID groups. For RAID group information, see [getrg on page 95](#).

If you do not specify the options, the system displays all the thin pool properties.

Note

If a storage pool has tiers with different RAID types, the pool RAID type is displayed as 'Mixed'.

SYNTAX

```
storagepool -list [-id poolID|-name poolName][-availableCap]
[-consumedCap] [-currentOp] [-description] [-disks] [-diskType]
[-luns] [-opState] [-opStatus] [-prcntOp] [-rawCap] [-rtype]
[-prcntFullThreshold] [-state] [-status] [-subscribedCap]
[-userCap] [-prcntFull] [-autoTiering] [-tiers[-rdrivecount]
[-loadBalance]] [-rebalance] [-fastcache]
[-lunAllocation] [-snapshotAllocation]
[-metadataAllocation] [-lunSubscribedCap]
[-snapshotSubscribedCap] [-metadataSubscribedCap]
[-compressionSavings] [-capacities]
[-snapPoolFullThresholdEnabled] [-snapPoolFullHWM]
[-snapPoolFullLWM] [-snapPoolFullState]
[-snapSpaceUsedThresholdEnabled] [-snapSpaceUsedHWM]
[-snapSpaceUsedLWM] [-snapSpaceUsedState]
[-dedupState] [-dedupRate] [-dedupTieringPolicy]
[-dedupInitialTier] [-efficiencySavings] [-dedupSharedCap] [-
all]
```

OPTIONS

`-id poolID|-name poolName`

Displays the name or ID of the thin pool with the specified name or ID. You cannot use the options `-id` and `-name` together.

`-availableCap`

Displays the free shared capacity (in blocks and GB).

`-consumedCap`

Displays the total allocated capacity in the shared capacity of the thin pool. The consumed capacity of a thin pool includes storage used by MLU for maintaining metadata (mapping info).

`-currentOp`

Displays the current operation.

-description
Returns the description of the thin pool.

-disks
Lists the disks used in each thin pool on the system. The disks are listed in *Bus_Enclosure_BankSlot* format, with each disk on a new line.

-diskType
Displays the disk type of the components in the pool. If the components have different types, the system displays Mixed.

-luns
Lists the IDs of all thin LUNs in the thin pool.

-opState
Displays the current operation state. The system displays N/A, if the current operation state cannot be determined.

-opStatus
Displays the additional descriptive information for the current state of the thin pool.

-prcntOp
Displays the additional descriptive information for the current state of the thin pool operation in progress. If there is no operation, the system displays 0.

-rawCap
Displays the raw capacity of the thin pools.

-rtype
Lists the thin pool RAID type.

-prcntFullThreshold
Displays the user modifiable used space threshold, beyond which the system generates the alerts.

-state
Displays the state of the thin pool.

-status
Displays the information for the current state of the thin pool.

-subscribedCap
Displays the host-visible capacity of the thin pool.

-userCap
Displays the user-accessible capacity of the thin pool.

-prcntFull
Displays the used percentage.

-autoTiering
Displays the auto-tiering state (scheduled/manual)

-tiers
Displays the tier name, RAID type, capacity, and disk information for each tier. If auto-tiering is enabled, the data to be relocated up/down will also be displayed.

-rdrivecount
Specifies the RAID drive count used to create the internal RAID groups.

-loadBalance
Displays the load balance data, which is the data targeted within the tier.

-rebalance
Displays the rebalancing percent completed.

-fastcache
Lists the FAST Cache state of the storage pool.

-lunAllocation
Specifies the total amount of storage allocated within a pool to the data that has been written to all the primary LUNs in that pool.

-snapshotAllocation
Specifies the total amount of storage allocated within a pool to all the VNX snapshots created in that pool.

-metadataAllocation

- Specifies the total amount of storage allocated within a pool to store the metadata for the LUNs and VNX snapshots in that pool.
- lunSubscribedCap
 - Specifies the total amount of user data that can be written to all of the LUNs in that Pool.
- snapshotSubscribedCap
 - Specifies the total amount of space that all the LUN snapshots consume, as well as what they would consume if all the primary LUN data would be overwritten.
- metadataSubscribedCap
 - Specifies the total amount of metadata storage required to support writing to the entire contents of all of the LUNs in the pool.
- compressionSavings
 - Specifies the total amount of storage saved by using compression on the LUNs in the pool.
- capacities
 - Specifies all the related capacity of a pool such as, Raw Capacity, User Capacity, Consumed Capacity, LUN Allocation, Snapshot Allocation, Metadata Allocation, Available Capacity, Percent Full, Total Subscribed Capacity, LUN Subscribed Capacity, Snapshot Subscribed Capacity, Metadata Subscribed Capacity, Percent Subscribed, Oversubscribed by, and Compression Savings.
- snapPoolFullThresholdEnabled
 - Specifies if the system should check the pool full high water mark for auto delete.
- snapPoolFullHWM
 - Specifies the pool full high watermark of the specified storage pool(s).
- snapPoolFullLWM
 - Specifies the pool full low watermark of the specified storage pool(s).
- snapPoolFullState
 - Specifies the pool full auto delete state.
- snapSpaceUsedThresholdEnabled
 - Specifies if the system should check the used snapshot space for high water mark for auto delete.
- snapSpaceUsedHWM
 - Specifies the used snapshot space for high watermark of the specified storage pool(s).
- snapSpaceUsedLWM
 - Specifies the used snapshot space for low watermark of the specified storage pool(s).
- snapSpaceUsedState
 - Specifies the used snapshot space in auto delete state.
- dedupState
 - Displays Deduplication State, Deduplication Status, Deduplication Percent Completed, and Deduplication Remaining Size of the specified storage pool. Possible values of Deduplication State are Idle (No Deduplicated LUNs), Pending, Running, or Paused.
- dedupRate
 - Displays the deduplication rate of the specified storage pool. Possible values are Low, Medium, or High.
- dedupTieringPolicy
 - Displays the common tiering policy for all deduplicated LUNs in this storage pool. Possible values are No Movement, Auto Tier, Highest Available, or Lowest Available.
- dedupInitialTier
 - Displays the common initial tier for all deduplicated LUNs in this storage pool. Possible values are Optimize Pool, Highest Available, or Lowest Available.
- efficiencySavings
 - Displays space savings from deduplication and VNX Snapshots for the specified storage pool.

`-dedupSharedCap`

Displays the consumed capacity of all deduplicated LUNs in the storage pool.

`-all`

Displays all the storage pool information, including the tiers (always) and auto-tiering state (if the auto-tiering enabler has been installed).

EXAMPLE #1

List the RAID type for a storage pool with tiers that are RAID 5 and RAID 6

```
naviseccli -h riker-spa storagepool -list -id 0 -rtype
```

```
Pool Name: Pool 0
Pool ID: 0
Raid Type: Mixed
```

EXAMPLE #2

```
naviseccli -h riker-spa storagepool -list -id 0 -tiers -rdrivecount
```

EXAMPLE #3

List all deduplication properties.

```
naviseccli -h riker-spa storagepool -list -dedupState
-dedupRatededupTieringPolicy -dedupInitialTier -efficiencySavings
-dedupSharedCap
```

```
Pool Name: Pool 0
Pool ID: 0
Deduplication State: Idle (No Deduplicated LUNs)
Deduplication Status: OK(0x0)
Deduplication Percent Completed: 100
Deduplication Remaining Size (Blocks): 0
Deduplication Remaining Size (GBs): 0.000
Optimal Deduplicated LUN SP Owner: N/A
Deduplication Rate: Medium
Deduplicated LUNs' Tiering Policy: Auto Tier
Deduplicated LUNs' Initial Tier: Highest Available
Efficiency Savings (Blocks): 20986976
Efficiency Savings (GBs): 10.007
Deduplication Shared Capacity (Blocks): 209869760
Deduplication Shared Capacity (GBs): 100.07
```

EXAMPLE #4

List all capacity properties.

```
naviseccli -h riker-spa storagepool -list -capacities
```

```
Pool Name: Pool 0
Pool ID: 0
Raw Capacity (Blocks): 2814421510
Raw Capacity (GBs): 1342.021
User Capacity (Blocks): 2246234112
User Capacity (GBs): 1071.088
Consumed Capacity (Blocks): 98758656
Consumed Capacity (GBs): 47.092
Available Capacity (Blocks): 21474754565
Available Capacity (GBs): 1023.996
Percent Full: 4.397
LUN Allocation (Blocks): 31457280
LUN Allocation (GBs): 15.000
Snapshot Allocation (Blocks): N/A3
Snapshot Allocation (GBs): N/A
Metadata Allocation (Blocks): 509317121
Metadata Allocation (GBs): 24.286
```

```

Total Subscribed Capacity (Blocks): 3984588800
Total Subscribed Capacity (GBs): 1900.000
Percent Subscribed: 177.390
Oversubscribed by (Blocks): 1738354688
Oversubscribed by (GBs): 828.912
LUN Subscribed Capacity (Blocks): 3856662528
LUN Subscribed Capacity (GBs): 1839.000
Snapshot Subscribed Capacity (Blocks): N/A4
Snapshot Subscribed Capacity (GBs): N/A
Metadata Subscribed Capacity (Blocks): 1069392962
Metadata Subscribed Capacity (GBs): 50.993
Compression Savings (Blocks): 0
Compression Savings (GBs): 0.000
Efficiency Savings (Blocks): 20986976
Efficiency Savings (GBs): 10.007
Deduplication Shared Capacity (Blocks): 209869760
Deduplication Shared Capacity (GBs): 100.07

```

EXAMPLE #5

List all properties.

```
naviseccli -h riker-spa storagepool -list -all
```

```

Pool Name: Pool 0
Pool ID: 0
Raid Type: r_5
Percent Full Threshold: 70
Description:
Disk Type: Fibre Channel
State: Ready
Status: OK(0x0)
Current Operation: None
Current Operation State: N/A
Current Operation Status: N/A
Current Operation Percent Completed: 0
Raw Capacity (Blocks): 1401742320
Raw Capacity (GBs): 668.403
User Capacity (Blocks): 1111558400
User Capacity (GBs): 530.032
Consumed Capacity (Blocks): 29361920
Consumed Capacity (GBs): 14.001
Available Capacity (Blocks): 1082196480
Available Capacity (GBs): 516.031
Percent Full: 2.642
Total Subscribed Capacity (Blocks): 25167360
Total Subscribed Capacity (GBs): 12.001
Percent Subscribed: 2.264
Oversubscribed by (Blocks): 0
Oversubscribed by (GBs): 0.000
Auto-Tiering: Manual

Tier Name: FC
Raid Type: r_5
User Capacity (GBs): 530.03
Consumed Capacity (GBs): 4.00
Available Capacity (GBs): 526.03
Percent Subscribed: 0.75%
Data Targeted for Higher Tier (GBs): 0.00
Data Targeted for Lower Tier (GBs): 0.00
Disks (Type):
Bus 0 Enclosure 0 Disk 13 (Fibre Channel)
Bus 0 Enclosure 0 Disk 11 (Fibre Channel)
Bus 0 Enclosure 0 Disk 14 (Fibre Channel)
Bus 0 Enclosure 0 Disk 12 (Fibre Channel)
Bus 0 Enclosure 0 Disk 10 (Fibre Channel)

Disks:

```

```

Bus 0 Enclosure 0 Disk 11
Bus 0 Enclosure 0 Disk 13
Bus 0 Enclosure 0 Disk 10
Bus 0 Enclosure 0 Disk 12
Bus 0 Enclosure 0 Disk 14
LUNs: 1, 0
Flash Cache: Enabled
Deduplication State: Running
Deduplication Status: OK(0x0)
Deduplication Percent Completed: 70
Deduplication Remaining Size (Blocks): 12608078
Deduplication Remaining Size (GBs): 3.006
Optimal Deduplicated LUN SP Owner: A
Deduplication Rate: Medium
Deduplicated LUNs' Tiering Policy: Auto Tier
Deduplicated LUNs' Initial Tier: Highest Available
Efficiency Savings (Blocks): 20986976
Efficiency Savings (GBs): 10.007
Deduplication Shared Capacity (Blocks): 209869760
Deduplication Shared Capacity (GBs): 100.07

```

EXAMPLE #6

Default output (deduplication properties do not display).

```
naviseccli -h riker-spa storagepool -list
```

```

Pool Name: Pool 0
Pool ID: 0
Raid Type: r_5
Percent Full Threshold: 70
Description:
Disk Type: Fibre Channel
State: Ready
Status: OK(0x0)
Current Operation: None
Current Operation State: N/A
Current Operation Status: N/A
Current Operation Percent Completed: 0
Raw Capacity (Blocks): 1401742320
Raw Capacity (GBs): 668.403
User Capacity (Blocks): 1111558400
User Capacity (GBs): 530.032
Consumed Capacity (Blocks): 29361920
Consumed Capacity (GBs): 14.001
Available Capacity (Blocks): 1082196480
Available Capacity (GBs): 516.031
Percent Full: 2.642
Total Subscribed Capacity (Blocks): 25167360
Total Subscribed Capacity (GBs): 12.001
Percent Subscribed: 2.264
Oversubscribed by (Blocks): 0
Oversubscribed by (GBs): 0.000
Disks:
Bus 0 Enclosure 0 Disk 11
Bus 0 Enclosure 0 Disk 13
Bus 0 Enclosure 0 Disk 10
Bus 0 Enclosure 0 Disk 12
Bus 0 Enclosure 0 Disk 14
LUNs: 1, 0

```

storagepool -destroy

Destroys the thin pool.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -destroy` command deletes a thin pool. You must specify the thin pool ID or the thin pool name to be deleted from the system.

SYNTAX

```
storagepool -destroy -id poolID |-name poolName [-o]
```

OPTIONS

`-id poolID`|-name *poolName*

Specifies the thin pool ID or name to be destroyed. You cannot use the options `-id` and `-name` together.

`-o`

Executes the command without prompting for confirmation.

OUTPUT

None if the command succeeds; status or error information if the command fails.

storagepool -modify

Modifies the thin pool.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -modify` command lets you modify certain properties of an existing thin pool. You must specify at least one property of the thin pool to be modified.

SYNTAX

```
storagepool -modify -id poolID|-name poolName[-newName
newName] [-description description] [-fastcache on|off]
[-prcntFullThreshold threshold] [-autotiering scheduled|
manual]
[-snapPoolFullThresholdEnabled on|off] [-snapPoolFullHWM
poolFullHWM] [-snapPoolFullLWM poolFullLWM]
[-snapSpaceUsedThresholdEnabled on|off] [-snapSpaceUsedHWM
snapSpaceUsedHWM] [-snapSpaceUsedLWM snapSpaceUsedLWM]
[-dedupState Running|Paused] [-dedupRate Low|Medium|High]
[-dedupTieringPolicy noMovement|autoTier|highestAvailable|
lowestAvailable]
[-dedupInitialTier optimizePool|highestAvailable|
lowestAvailable] [-o]
[-startDedup force]
```

OPTIONS

- id *poolID* | -name *poolName*
Specifies the thin pool ID or name to be modified. You cannot use the options `-id` and `-name` together.
- newName
Sets the unique name for the thin pool. The character length ranges from 1 to 64.
- description
Sets the new description of the thin pool. The character length ranges from 0 to 255.
- fastcache on|off
Modifies the FAST Cache configuration of the storage pool.
- prcntFullThreshold
Specifies the new user capacity threshold, beyond which the system generates the alerts. The value ranges from 1 to 84. The default value is 70%.
- autotiering scheduled|manual
Sets the auto-tiering state to `scheduled` or `manual` for the storage pool.
- snapPoolFullThresholdEnabled
Indicates that the system should check the pool full high water mark for auto delete. Valid values are `On` and `Off`.
- snapPoolFullHWM
Specifies the new pool full high watermark that triggers the auto delete process of the storage pool to be modified.
- snapPoolFullLWM
Specifies the new pool full low watermark that stops the auto delete process of the storage pool to be modified.
- snapSpaceUsedThresholdEnabled
Indicates that the system should check the snapshot space used high water mark for auto delete. Valid values are `On` and `Off`.
- snapSpaceUsedHWM
Specifies the new snapshot space used high watermark that triggers the auto delete process of the storage pool to be modified.
- snapSpaceUsedLWM
Specifies the new snapshot space used low watermark that stops the auto delete process of the storage pool to be modified.
- dedupState
Allows you to pause or resume deduplication on the specified storage pool. Valid values are `Running` or `Paused`.
- dedupRate
Specifies the deduplication rate on the specified storage pool. Valid values are `Low`, `Medium`, or `High`.
- dedupTieringPolicy
Specifies the common FAST VP tiering policy for all deduplicated LUNs in the specified storage pool. Valid values are `noMovement`, `autoTier`, `highestAvailable`, or `lowestAvailable`.
- dedupInitialTier
Specifies the common FAST VP initial tier for all deduplicated LUNs in the specified storage pool. Valid values are `optimizePool`, `highestAvailable`, or `lowestAvailable`.
- o
Force the pool modification by ignoring any potential warnings.
- startDedup *force*
Starts a pool-level deduplication operation immediately. All restrictions are bypassed, including the less than 64G threshold, for one deduplication pass. If three deduplication operations are currently running, the new request will wait until one session finishes before starting.

Note

This option should be used on a limited basis. Forcing a deduplication operation could affect future deduplication jobs.

EXAMPLE #1

```
naviseccli -h riker-spa storagepool -modify -poolid 0 -dedupRate high
-dedupTieringPolicy autoTier -dedupState paused -dedupInitialTier
highestAvailable
```

EXAMPLE #2

```
naviseccli -h <array_ip> storagepool -modify -poolid 0 -startDedup
force
```

```
You are about to override the default deduplication schedule for this
pool. Forcing deduplication can impact future deduplication
opportunities.
```

```
Are you sure you want to perform this operation?(y/n):
```

storagepool -expand

Expands thin pool capacity.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -expand` command lets you expand the physical capacity of the thin pool on the system. If you need to expand an existing tier of the selected set of disks, the current RAID type of the tier is used to expand the tier. However, if you specify a RAID type that doesn't match the RAID type of the existing tier, the system displays an error message.

SYNTAX

```
storagepool -expand -id poolID | -name poolName -disks
disksList [-rtype raidType [-rdrivecount drivecount]] [-
initialverify yes|no] [-skipRules] [-o]
```

OPTIONS

`-id poolID | -name poolName`

Specifies the thin pool ID or name to be expanded. You cannot use the options `-id` and `-name` together.

`-rtype`

Specifies the type of RAID group of the thin pool to be expanded. The valid RAID groups are `r_5`, `r_6`, and `r_10`.

`-disks`

Specifies the list of disks in `Bus_Enclosure_BankSlot` notation. If you do not specify bus in the format, Bus 0 is assumed.

`-o`

Executes the command without prompting for confirmation.

`-initialverify`

Allows the user to specify not to have initial background verify during pool creation and expansion in order to improve operation performance of pool creation and expansion.

`-skipRules`

Skips the best practices check while expanding thin pools. If you do not specify this option, the thin pool expansion may fail due to any error checked in the thin rules. The default value is `FALSE`.

EXAMPLE #1

```
naviseccli -h riker-spa storagepool -modify -id 0 -disks 0_1_14
0_2_0 ... 0_2_14 -rtype r_6 -rdrivecount 16
```

EXAMPLE #2

Expanding a storage pool with `NL_SAS` drives using RAID 6. The storage pool does not have any disks in the Capacity tier, therefore the tier is created using RAID 6

```
naviseccli -h riker-spa storagepool -expand -id 0 -rtype r_6 -disks
0_1_0 0_1_1 0_1_2 0_1_3 0_1_4 0_1_5 0_1_6 0_1_7 -o
```

OUTPUT

None if the command succeeds; status or error information if the command fails.

storagepool -cancelexpand

Cancels a failed expansion.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -cancelexpand` command lets you cancel a failed expansion on a storage pool. You must specify the thin pool ID or thin pool name to cancel the expand operation.

SYNTAX

```
storagepool -cancelexpand -id poolID| -name poolName
```

OPTIONS

`-id poolID`| `-name poolName`

Specifies the thin pool ID or name. You cannot use the options `-id` and `-name` together.

OUTPUT

None if the command succeeds; status or error information if the command fails.

storagepool -feature -info

Lists the pool configuration information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `storagepool -feature -info` command displays the pool configuration information. If the options are not specified, then all the values are displayed.

SYNTAX

```
storagepool -feature -info [-isVirtualProvisioningSupported]
[-maxPools] [-maxDiskDrivesPerPool] [-maxDiskDrivesAllPools]
[-maxDiskDrivesPerOp] [-maxPoolLUNs] [-minPoolLUNSize]
[-maxPoolLUNSize] [-numPools] [-numPoolLUNs] [-numThinLUNs]
[numDiskDrivesAllPools] [-availableDisks]
```

OPTIONS

`-isVirtualProvisioningSupported`
Specifies whether the thin enabler is installed on the system.

`-maxPools`
Displays the maximum number of the storage pools that can be created on the system.

`-maxDiskDrivesPerPool`
Specifies the maximum number of the disk drives in a pool.

`-maxDiskDrivesAllPools`
Specifies the maximum number of the disk drives in all pools.

`-maxDiskDrivesPerOp`
Specifies the maximum number of the disk drives used per operation.

`-maxPoolLUNs`
Displays the maximum number of storage pool LUNs that can be created on the system.

`-minPoolLUNSize`
Displays the minimum pool LUN size supported (as reported by the MLU driver) in blocks and MBs.

`-maxPoolLUNSize`
Displays the maximum pool LUN size supported (as reported by the MLU driver) in blocks and TBs.

`-numPools`
Displays the current number of storage pools created on the system.

`-numPoolLUNs`
Displays the current number of storage pool LUNs created on the system.

`-numThinLUNs`
Displays the total number of all pool LUNs created that are thin and non-thin.

`numDiskDrivesAllPools`
Displays the number of disk drives used in all pools.

`-availableDisks`
A line feed separated list of disks that are available for use by the Virtual Pools feature.

OUTPUT

```
Is Virtual Provisioning Supported: true
Max. Pools: 120
Max. Disks Per Pool: 120
Max. Disks for all Pools: 0
Max. Disks per Operation: 120
Max. Pool LUNs: 4094
Min. Pool LUN Size(Blocks): 2097152
Max. Pool LUN Size(Blocks): 30064771072
Max. Pool LUN Size(GBs): 14336.000
Total Number of Pools: 1
Total Number of Pool LUNs: 0
Total Number of all Pool LUNs that are thin: 20
Total Number of all Pool LUNs that are non-thin: 20
Number of Disks used in Pools: 0
Available Disks:
Bus 1 Enclosure 1 Disk 9
```



```

Bus 1 Enclosure 1 Disk 8
Bus 1 Enclosure 1 Disk 7
Bus 1 Enclosure 1 Disk 6
Bus 1 Enclosure 1 Disk 5
Bus 1 Enclosure 1 Disk 4
Bus 1 Enclosure 1 Disk 3
Bus 1 Enclosure 1 Disk 2
Bus 1 Enclosure 1 Disk 1
Bus 1 Enclosure 1 Disk 0
Bus 1 Enclosure 0 Disk 9
Bus 1 Enclosure 0 Disk 8
Bus 1 Enclosure 0 Disk 7
Bus 1 Enclosure 0 Disk 6
Bus 1 Enclosure 0 Disk 5
Bus 1 Enclosure 0 Disk 4
Bus 1 Enclosure 0 Disk 3
Bus 1 Enclosure 0 Disk 2
Bus 1 Enclosure 0 Disk 1
Bus 1 Enclosure 0 Disk 0
Bus 0 Enclosure 0 Disk 14
Bus 0 Enclosure 0 Disk 13
Bus 0 Enclosure 0 Disk 12
Bus 0 Enclosure 0 Disk 11
Bus 0 Enclosure 0 Disk 10

```

storagepool -resumeharvesting

Resumes harvesting on an existing storage pool.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli storagepool -resumeharvesting` command lets you resume harvesting on an existing storage pool.

SYNTAX

```
storagepool -resumeharvesting -id poolID|-name poolName
```

OPTIONS

```
-id poolID|-name poolName
```

Specifies the ID or name of the storage pool on which the resume harvesting operation takes place.

OUTPUT

None if the command succeeds; status or error information if the command fails.

lun -create

Creates a LUN.

PREREQUISITES

This command is available only on subsystems running Unisphere version 6.30 and later.

To create thin LUNs, you must have the Thin Provisioning Enabler installed on the subsystem.

DESCRIPTION

The `naviseccli lun -create` command lets you create pool LUNs or thin LUNs in the storage pool.

SYNTAX

```

lun -create [-type Thin|nonThin]
-capacity capacity [-sq mb|gb|tb|bc]
-poolId storagePoolID |-poolname storagePoolName -sp SPID
[-aa autoAssignment]
[-l lunNumber]
[-name lunName]
[-offset logicalBlockAddress]
[-tieringPolicy noMovement|autoTier|highestAvailable|
lowestAvailable]
[-initialTier optimizePool|highestAvailable|lowestAvailable]
[-deduplication On|Off]

```

OPTIONS

-type Thin|nonThin

Specifies whether a thin LUN or pool LUN (non-thin LUN) is created. If not specified and `poolID` or `poolName` is provided, a pool LUN will be created.

Note

To create a thin pool LUN, you must explicitly use the `-type Thin` option, or a thick pool LUN is created.

-capacity *capacity*

Specifies the storage capacity of the LUN.

-sq mb|gb|tb|bc

Specifies the size qualifier for the capacity. The valid values are gb, tb, and bc. The default value is GB.

-poolId *storagePoolID*|-poolName *storagePoolName*

Specifies the storage pool ID or storage pool name in which the new LUN will be created.

-sp *SPID*

Specifies the default SP to which the new LUN must belong. The valid values are A and B.

-aa *autoAssignment*

Specifies whether auto assignment is enabled by default or not. The valid values are 1 (yes) and 0 (no). The default value is 0.

-l *lunNumber*

Specifies the LUN number. The lower limit is 0 and the upper limit is the number of LUNs supported on the platform. The default value is the next lowest available value.

-name *lunName*

Specifies the LUN name. The name has to be unique amongst the pool LUNs. If you do not specify a name, the system defaults to LUN <id>.

-offset *logicalBlockAddress*

Sets the LUN's starting Logical Block Address (LBA) to begin at a specified offset location. The valid value is any unsigned decimal number. The default value is 0.

-tieringPolicy noMovement|autoTier|highestAvailable|lowestAvailable

Sets the auto-tiering policy. If the auto-tiering policy is not set when creating the LUN, the default policy will depend on the `initialTier` settings. The auto-tiering policy can be set to:

noMovement

No data relocation between tiers but the initial tier is used to guide initial placement of data.

autoTier

Slices are subject to relocation based on the stats collected. This setting requires that the initial tier also be set to optimize pool.

highestAvailable

Slices are subject to relocation with a bias to the highest tier. This setting requires that the initial tier also be set to the highest available tier.

lowestAvailable

Slices are subject to relocation with a bias to the lowest tier. This setting requires that the initial tier also be set to the lowest available tier.

`-initialTier optimizePool|highestAvailable|lowestAvailable`

Sets the initial tier preference.

optimizePool

No preference is specified and the default data placement algorithm applies.

highestAvailable

Sets the preferred tier for initial data placement to be the highest tier.

lowestAvailable

Sets the preferred tier for initial data placement to be the lowest tier.

Note

The following combination is now a valid combination and does not return an incompatible options error: `lun -create [...] -tieringPolicy autoTier -initialTier highestAvailable`.

The combination of `tieringPolicy` and `initialTier` is now compatible for the Start High then Auto-tier option.

`-deduplication`

Specifies whether deduplication is enabled on this LUN. Valid values are On or Off.

EXAMPLE

```
naviseccli -h riker-spa lun -create -deduplication on
```

OUTPUT

Returns 0 if the command succeeds; status or error information if it fails.

lun -modify

Modifies the LUN properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli lun -modify` command lets you modify the properties of an existing LUN. At least one property to modify a LUN should be specified to issue this command. The `lun -modify` command provides the same options for all pool LUNs.

SYNTAX

```
lun -modify -l lunNumber|-name lunName
[-aa Autoassign]
```

```

[-newname lunName]
[-sp newDefaultSPID]
[-tieringPolicy noMovement|autoTier|highestAvailable|
lowestAvailable]
[-initialTier
optimizePool|highestAvailable|lowestAvailable]
[-deduplication On|Off [-destroySnapshots]] [-o]

```

OPTIONS

- l *lunNumber*
Specifies the LUN number. The valid value is a number that represents an existing LUN.
- name *lunName*
Specifies the LUN name. The name has to be unique among the LUNs in the storage pool.
- aa *Autoassign*
Specifies whether auto assignment is enabled by default or not. The valid values are 1 (yes) and 0 (no).
- newname *lunName*
Specifies the new LUN name. The name has to be unique among the LUNs in the storage pool.
- sp *newDefaultSPID*
Specifies the default SP to which the LUN belongs. The valid values are A and B. If the value is changed from the allocation owner, a prompt displays a warning about potential performance impact. Select Yes to proceed.
- tieringPolicy noMovement|autoTier|highestAvailable|lowestAvailable
Modifies the auto-tiering policy. The auto-tiering policy can be set to:
 - noMovement**
No data relocation between tiers but the initial tier is used to guide initial placement of data.
 - autoTier**
Slices are subject to relocation based on stats collected. This setting requires that the initial tier also be set to optimize pool.
 - highestAvailable**
Slices are subject to relocation with a bias to the highest tier. This setting requires that the initial tier also be set to the highest available tier.
 - lowestAvailable**
Slices are subject to relocation with a bias to the lowest tier. This setting requires that the initial tier also be set to the lowest available tier.
- initialTier optimizePool|highestAvailable|lowestAvailable
Sets the initial tier preference.
 - optimizePool**
No preference is specified and the default data placement algorithm applies.
 - highestAvailable**
Sets the preferred tier for initial data placement to be the highest tier.
 - lowestAvailable**
Sets the preferred tier for initial data placement to be the lowest tier.

Note

The combinations `lun -modify [...] -tieringPolicy autoTier -initialTier highestAvailable` and `lun -modify [...] -initialTier highestAvailable` are now a valid combination and do not return an incompatible options error. These combinations are now compatible for the Start High then Auto-tier option.

`-deduplication On|Off`

Specifies whether deduplication is enabled on this LUN. Valid values are On or Off.

Note

You cannot modify deduplication settings for a LUN that belongs to file storage. To modify deduplication settings for a file storage pool or a file system, use Unisphere or the VNX for File CLI

`-destroySnapshots`

Specifies to explicitly destroy associated VNX Snapshots before turning deduplication on or off for the specified LUN. If a LUN has associated VNX Snapshots and you do not use the `-destroySnapshots` option, you cannot change the deduplication value.

`-o`

Overrides the confirmation message.

EXAMPLE

```
naviseccli lun -modify -l LUN3 -deduplication on
```

```
"SP A" is the optimal owner for the deduplicated LUNs in the storage pool. Enabling deduplication on "LUN3" owned by "SP B" may greatly impact the performance of deduplication. It is recommended that you change the default owner to avoid performance impact.
```

```
Are you sure you want to perform this operation? (y/n):
```

lun -expand

Expands a LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli lun -expand` lets you expand a pool LUN. This command will not work for classic LUNs. If the new capacity specified is higher than the existing capacity and the requirements for expansion are met, the LUN would expand. If the new capacity is less than the existing capacity, an error message will be displayed.

SYNTAX

```
lun -expand -l lunNumber|-name lunName-capacity capacity
[-sq mb|gb|tb|bc] [-o]
```

OPTIONS

`-l lunNumber`

Specifies the LUN number. The valid value is a number that represents an existing LUN.

`-name lunName`
 Specifies the LUN name. The name has to be unique among the pool LUNs.

`-capacity capacity`
 Specifies the new capacity of the LUN. The new capacity should be higher than the existing capacity for LUN expansion.

`-sq mb|gb|tb|bc`
 Specifies the size qualifier for the capacity. The valid values are mb, gb, tb, and bc. The default value is GB.

OUTPUT

Returns 0 if the command succeeds; status or error information if it fails.

lun -destroy

Destroys a LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli lun -destroy` command lets you destroy a LUN. The LUN number or LUN name is used to destroy any type of LUN since they are unique across all LUNs in the storage pools.

SYNTAX

```
lun -destroy -l lunNumber|-name lunName [-o]
```

OPTIONS

`-l lunNumber`
 Specifies the LUN number.

`-name lunName`
 Specifies the LUN name.

OUTPUT

Returns 0 if the command succeeds; status or error information if it fails.

lun -list

Lists the LUN properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli lun -list` lets you list all the LUN properties. If the options are not specified, then all the values will be returned for all the LUNs. If only one option is specified, it will be displayed along with the Logical Unit Number and Name.

SYNTAX

```
lun -list [-l lunNumber|-name lunName] | [-showOnly Thin|
nonThin|Compressed|nonCompressed|Deduped|NonDeduped]
[-primaryLunName primaryLunName]
[-aa] [-alowner] [-at] [-userCap] [-consumedCap] [-default]
[-drivetype] [-offset] [-rtype] [-state] [-status] [-owner]
[-poolName] [-uid] [-opDetails] [-isPoolLun] [-isThinLUN]
[-isCompressed] [-isPrivate] [-featureData] [-perfData]
```

```
[-tieringPolicy] [-initialTier] [-tiers]
[-dedupState] [-all]
```

OPTIONS

- l *lunNumber*
Specifies the LUN number in unsigned decimal format. If the LUN number is not found, an error message will be returned indicating the LUN may not exist.
- name *lunName*
Specifies the LUN name. If the LUN name is not found, an error message will be returned indicating the LUN may not exist.
- showOnly Thin|nonThin|Compressed|nonCompressed|Deduped|NonDeduped
Specifies the types of LUNs to be listed. The value Thin lists all thin LUNs, nonThin lists all non-thin LUNs, Compressed shows all compressed LUNs, nonCompressed shows all non-compressed LUNs, Deduped shows only deduplicated LUNs, and NonDeduped shows only non-deduplicated LUNs. The options Compressed and NonCompressed will be shown in the command usage even if the compression feature is not enabled.
- primaryLun
Show only the Snapshot Mount Point associated with the specified primary LUN number.
- primaryLunName
Show only the Snapshot Mount Point associated with the specified primary LUN name.
- aa
Specifies the auto assign property.
- alowner
Specifies the allocation owner of the LUN.
- at
Specifies the auto-trespass property.
- userCap
Specifies the user (host visible) capacity of the LUN in blocks and GBs.
- consumedCap
Specifies the consumed capacity of the LUN. The consumed capacity is always larger than the user capacity for pool LUNs.
- default
Specifies the default owner of the LUN.
- drivetype
Displays the drive type of the LUN which is essentially the drive type of the storage pool. If more than one type of drive are present in the pool, then Mixed will be displayed.
- offset
Displays the alignment offset of the LUN that was assigned when the LUN was bound.
- opDetails
Displays the details of the currently running operation. The values are name, state, status, and percent completed. If no current operation is running, then the values for these fields are N/A, N/A, N/A, and 0, respectively. This could be used to show compression and auto-tiering related information.
- owner
Displays the name of the SP that currently owns the LUN. If neither SP owns the LUN a value of Unknown is returned.
- isPoolLUN
Displays Yes, if it is either a thin LUN or pool LUN.
- isThinLUN
Displays Yes if it is a thin LUN.
- isCompressed

- Displays the isCompressed property of the LUN.
- poolName
Displays the name of the storage pool to which the mapped LUN belongs.
- rtype
Specifies the RAID group type of the LUN.
- perfData
Specifies the following information: Statistics Logging Current Time, Read Requests, Write Requests, Blocks Read, Blocks Written, Busy Ticks, Idle Ticks, Sum of Outstanding Requests, Non-Zero Request Count Arrivals, Implicit Trespasses, and Explicit Trespasses.
- isPrivate
Specifies if the LUN is a private LUN. Pool LUNs can be private since they can be part of a reserved LUN pool.
- featureData
Returns the information about the features on the LUN.
- state
Displays the state of the mapped LUN.
- status
Displays the status of the mapped LUN. This displays a string that corresponds to the status code, along with the status code in hex in lowercase, enclosed in brackets.
- uid
Displays the unique identifier (System WWN + LUN WWN) for the LUN.
- tieringPolicy
Displays the auto-tiering policy.
- initialTier
Displays the initial tier information.
- tiers
Displays the tier information.
- dedupState
Displays the Deduplication State and Deduplication Status. Possible values of Deduplication State are Enabling, On, Disabling, or Off.
- all
Displays the LUN information along with the auto-tiering policy (if the auto-tiering enabler is installed), the initial tier, and the tier information.

EXAMPLE #1

LUN has deduplication enabled.

```
naviseccli -h riker-spa lun -list -dedupState
```

```
LOGICAL UNIT NUMBER 0
Name: LUN 0
Deduplication State: Enabling(80% complete)
Deduplication Status: OK(0x0)
```

EXAMPLE #2

List only deduplication-enabled LUNs.

```
naviseccli -h riker-spa lun -list -showOnly Deduped -dedupState
```

```
LOGICAL UNIT NUMBER 0
Name: LUN 0
Deduplication State: On
Deduplication Status: OK(0x0)
```

EXAMPLE #3

List all properties.


```
naviseccli -h riker-spa lun -list -all
```

```
LOGICAL UNIT NUMBER
LOGICAL UNIT NUMBER 0
Name: LUN 0
UID: 60:06:01:60:83:32:28:00:3A:23:E6:E2:5C:37:DF:11
Current Owner: SP A
Default Owner: SP A
Allocation Owner: SP A
User Capacity (Blocks): 209715200
User Capacity (GBs): 100.000
Consumed Capacity (Blocks): 218117120
Consumed Capacity (GBs): 104.006
Pool Name: Pool 0
Raid Type: r_5
Disk Type: Mixed
Offset: 0
Auto-Assign Enabled: DISABLED
Auto-Trespass Enabled: DISABLED
Current State: Ready
Status: OK(0x0)
Is Faulted: false
Is Transitioning: false
Current Operation: None
Current Operation State: N/A
Current Operation Status: N/A
Current Operation Percent Completed: 0
Statistics Logging Current Time: 04/06/10 08:07:09
Read Requests: 0
Read Requests SP A: 0
Read Requests SP B: 0
Write Requests: 0
Write Requests SP A: 0
Write Requests SP B: 0
Blocks Read: 0
Blocks Read SP A: 0
Blocks Read SP B: 0
Blocks Written: 0
Blocks Written SP A: 0
Blocks Written SP B: 0
Busy Ticks: 0
Busy Ticks SP A: 0
Busy Ticks SP B: 0
Idle Ticks: 0
Idle Ticks SP A: 0
Idle Ticks SP B: 0
Sum of Outstanding Requests: 0
Sum of Outstanding Requests SP A: 0
Sum of Outstanding Requests SP B: 0
Non-Zero Request Count Arrivals: 0
Non-Zero Request Count Arrivals SP A: 0
Non-Zero Request Count Arrivals SP B: 0
Implicit Trespasses: 0
Implicit Trespasses SP A: 0
Implicit Trespasses SP B: 0
Explicit Trespasses: 0
Explicit Trespasses SP A: 0
Explicit Trespasses SP B: 0
Is Pool LUN: Yes
Is Thin LUN: No
Is Private: No
Is Compressed: No
Features: N/A
Tiering Policy: Lowest Available
Initial Tier: Lowest Available
Tier Distribution:
FC: 100.00%
```

EXAMPLE #4

Default output (deduplication properties do not display).

```
naviseccli -h riker-spa lun -list
```

```
LOGICAL UNIT NUMBER 0
Name: LUN 0
```

deduplication -feature -pause

Pauses deduplication on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli deduplication -feature -pause` command pauses all underlying deduplication processes on the system. You must manually resume deduplication to cancel a pause. The system does not cancel the pause automatically.

SYNTAX

```
deduplication -feature -pause
```

EXAMPLE

```
naviseccli deduplication -feature -pause
```

OUTPUT

None if the command succeeds; status or error information if it fails.

deduplication -feature -resume

Resumes deduplication on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli deduplication -feature -resume` command resumes all underlying deduplication processes on the system.

SYNTAX

```
deduplication -feature -resume
```

EXAMPLE

```
naviseccli deduplication -feature -resume
```

OUTPUT

None if the command succeeds; status or error information if it fails.

deduplication -feature -info

Displays the deduplication feature information.

DESCRIPTION

The `naviseccli deduplication -feature -info` command displays the deduplication feature information.

SYNTAX

```
deduplication -feature -info
```

EXAMPLE # 1

Deduplication feature is enabled.

```
naviseccli -h <array_ip> deduplication -feature -info
```

```
Is Deduplication Supported: true  
Deduplication Feature State: On
```

EXAMPLE # 2

Deduplication feature is paused.

```
naviseccli -h <array_ip> deduplication -feature -info
```

```
Is Deduplication Supported: true  
Deduplication Feature State: Paused
```


CHAPTER 4

Auto-tiering Commands

This chapter provides a brief overview on the auto-tiering feature and includes a list of commands that you can use if the feature is enabled on your system.

Major topics include:

- [About tiered storage](#)..... 334
- [autotiering -schedule -modify](#)..... 335
- [autotiering -schedule -enable|disable](#)..... 336
- [autotiering -relocation -start](#)..... 336
- [autotiering -relocation -setrate](#)..... 337
- [autotiering -relocation -pause|resume](#)..... 338
- [autotiering -relocation -stop](#)..... 338
- [autotiering -info](#)..... 339

About tiered storage

Tiered storage is the assignment of different categories of data to different types of storage media to reduce the total storage cost. Data categories may be based on the levels of protection needed, performance requirements, frequency of use, cost, and other considerations. The purpose of tiered storage is to retain the most frequently accessed or important data on fast, high performance (more expensive) drives, and move the less frequently accessed and less important data to low performance (less expensive) drives. The software provides two types of tiered storage:

- Initial tier placement (does not require the auto-tiering enabler)
- Auto-tiering placement (requires the auto-tiering enabler)

Initial tier placement (auto-tiering enabler is not installed)

For initial tier placement, you must manually specify the storage tier to place the LUN data. The tier choices are **Optimize for Pool Performance**, **Highest Available Tier**, and **Lowest Available Tier**. Select the tier based on the relative performance requirements for the LUN data. After you assign a LUN to a specific tier, you must perform a LUN migration to relocate the data to a different tier or install the auto-tiering enabler.

Table 19 Initial tier placement (auto-tiering enabler is not installed)

Initial tier placement	Description
Optimize for Pool Performance	No tier preference is specified. The system uses the default data placement algorithm to place the data.
Highest Available Tier	Sets the preferred tier for initial data placement to the highest available tier.
Lowest Available Tier	Sets the preferred tier for initial data placement to the lowest available tier.

Auto-tiering placement (auto-tiering enabler is installed)

The auto-tiering feature migrates data between storage tiers to lower the customer's total cost of ownership (TCO). Storage pools are configured with multiple classes of storage devices (EFD, FC, and SATA) and system software continually tracks the usage of the data stored on the LUNs. Using these LUN statistics, the auto-tiering feature relocates data blocks (slices) of each LUN to the storage tier that is best suited for that data.

Table 20 Auto-tiering placement (auto-tiering enabler is installed)

Auto-tiering placement	Description
Auto Tier	The system uses the LUN performance statistics for data relocation.
Highest Available Tier	Sets the preferred tier for data relocation to the highest available tier.
Lowest Available Tier	Sets the preferred tier for data relocation to the lowest available tier.
No Data Movement	No data relocation between tiers. The initial tier placement is retained.

If you install the auto-tiering enabler on a system with initial tier placement settings, the system bases the auto-tier settings on the initial tier settings.

Table 21 Initial tier plus auto-tiering placement settings

Initial tier placement	Auto-tiering placement
Optimize for Pool Performance	Auto Tier
Highest Available Tier	Highest Available Tier
Lowest Available Tier	Lowest Available Tier
None	No Data Movement

naviseccli auto-tiering commands

Table 22 naviseccli auto-tiering commands

Command	Description
<code>autotiering -schedule -modify</code>	Modifies the schedule.
<code>autotiering -schedule -enable -disable</code>	Enables or disables the default schedule.
<code>autotiering -relocation -start</code>	Starts a relocation.
<code>autotiering -relocation -setrate</code>	Sets the relocation rate.
<code>autotiering -relocation -pause -resume</code>	Pauses or resumes relocations at the system level.
<code>autotiering -relocation -stop</code>	Stops a relocation.
<code>autotiering -info</code>	Displays the auto-tiering information.

autotiering -schedule -modify

Modifies the schedule.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -schedule -modify` command changes a scheduled day, start time, and duration. When modifying the schedule, the confirmation message contains the scheduled stop time. The stop time is calculated on the system by using the client's GMT offset, the scheduled start time, and scheduled duration.

SYNTAX

```
autotiering -schedule -modify [-day sun|mon|tue|wed|thu|fri|
sat]
[-starttime hh:mm]
[-duration hh:mm]
[-o]
```

OPTIONS

- day sun|mon|tue|wed|thu|fri|sat
Specified as a space separated list containing one or more of the following abbreviations: Sun, Mon, Tue, Wed, Thu, Fri, or Sat.
- starttime *hh:mm*
Specifies the scheduled start time. The hours are specified from 0-23 and the minutes are specified from 00-59. The start time is set relevant to the client's GMT offset.
- duration *hh:mm*
Specifies the scheduled duration. The maximum duration is 23 hours and 59 minutes. The minimum duration is 1 minute.
- o
Executes the command without prompting for the confirmation.

EXAMPLE

```
autotiering -schedule -modify -day Mon Wed Fri -starttime 10:30
-duration 2:00
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -schedule -enable|-disable

Enables or disables the default schedule.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -schedule -disable` command disables the auto-tiering default schedule. The default schedule can be enabled using the `schedule -enable` command.

SYNTAX

```
autotiering -schedule -enable|-disable [-o]
```

OPTIONS

- o
Executes the command without prompting for the confirmation.

EXAMPLE #1

```
autotiering -schedule -enable -o
```

EXAMPLE #2

```
autotiering -schedule -disable -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -relocation -start

Starts a relocation.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -relocation -start` command manually starts a relocation on a storage pool regardless of the auto-tiering state.

SYNTAX

```
autotiering -relocation -start -poolId poolId1 [poolId2] |
-poolname poolname] [poolname2] | -all
[-duration hh:mm]
[-rate high|medium|low]
[-o]
```

OPTIONS

- `-poolID poolID` `-poolname poolname`
Specifies the storage pool ID or the storage pool name on which the manual relocation starts.
- `-all`
Starts the manual relocation on all the storage pools on the system.
- `-duration hh:mm`
Specifies the duration for the manual relocation. The maximum duration is 23 hours and 59 minutes and the minimum duration is 1 minute. The default duration is 8 hours.
- `-rate high|medium|low`
Specifies the rate for a manual relocation. The valid values are high, medium, or low. If the rate is not specified, it defaults to the system rate, that is, the value returned from the `autotiering -info -rate`.
- `-o`
Executes the command without prompting for the confirmation.

EXAMPLE

```
autotiering -relocation -start -all -rate high -duration 10:00
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -relocation -setrate

Sets the relocation rate.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -relocation -setrate` command sets the relocation rate that determines how the scheduled relocations will be performed on the system. If the rate is not specified when starting a manual relocation, it will use this rate.

SYNTAX

```
autotiering -relocation -setrate -rate high|medium|low
[-o]
```

OPTIONS

- `-rate high|medium|low`
Sets the scheduled relocation rate. The valid values are high, medium, or low.
- `-o`
Executes the command without prompting for the confirmation.

EXAMPLE

```
autotiering -relocation -setrate -rate low -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -relocation -pause|-resume

Pauses or resumes relocations at the system level.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -relocation -pause` command pauses all the active relocation tasks on the system. It also prevents running new relocations (manual and scheduled) by immediately placing the tasks in the paused state.

The `autotiering -relocation -resume` command allows relocation tasks in the paused state to continue and allows new relocations (manual and scheduled) to start immediately.

SYNTAX

```
autotiering -relocation -pause|-resume [-o]
```

OPTIONS

-o

Executes the command without prompting for the confirmation.

EXAMPLE #1

```
autotiering -relocation -pause -o
```

EXAMPLE #2

```
autotiering -relocation -resume
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -relocation -stop

Stops a relocation.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -relocation -stop` command stops a relocation that was started manually or by the scheduler.

SYNTAX

```
autoTiering -relocation -stop -poolId poolID1 [poolID2] |  
-poolname pool name1 [pool name2] | -all [-o]
```

OPTIONS

-poolId *poolid*|-poolname *poolname*

Specifies the storage pool ID or the storage pool name on which the manual relocation stops.

- all
Stops the manual relocation on all the storage pools on the system.
- o
Executes the command without prompting for the confirmation.

EXAMPLE

```
autotiering -relocation -stop -all
```

OUTPUT

None if the command succeeds; status or error information if it fails.

autotiering -info

Displays the auto-tiering information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `autotiering -info` command displays the auto-tiering state, schedules, relocation rate, and operation status.

SYNTAX

```
autotiering -info [-state]
[-schedule]
[-rate]
[-opstatus [-loadBalance] [-poolId-poolId1 [-poolId 2...]] |
-poolName -poolName1 [-poolName2...]]]
```

OPTIONS

- state
Displays the auto-tiering state.
- schedule
Displays the auto-tiering schedule.
- rate
Displays the auto-tiering relocation rate.
- opstatus
Displays the auto-tiering operations status for all the storage pools on the system or for a given list of storage pools. If a relocation is not in progress for a storage pool, the status of the most recent relocation will be displayed.
- loadBalance
Displays the data to move within a tier. This data is also called the Load Balancing data.
- poolId *poolID1* [*poolID2...*] [-poolName *poolName1* [*poolName2...*]]
Specifies the storage pool ID or the storage pool name.

OUTPUT

```
naviseccli -h 10.5.2.182 autotiering -info -opstatus
```

The data to move within tiers information is displayed only if the `-loadBalance` option is used.

```
Storage Pool Name: Pool 0
Storage Pool ID: 0
Relocation Start Time: 02/01/13 10:07
Relocation Stop Time: N/A
Relocation Status: Active
Relocation Type: Rebalance
```

```
Relocation Rate: Low  
Data to Move Up (GBs): 1.25  
Data to Move Down (GBs): 3.51  
Data to Move Within Tiers (GBs): 1.75  
Data Movement Completed (GBs): 6.01  
Estimated Time to Complete: 2 minutes  
Schedule Duration Remaining: N/A
```

Output for a rebalance in progress:

```
Storage Pool Name: Pool 0  
Storage Pool ID: 0  
Relocation Start Time: 02/01/13 10:07  
Relocation Stop Time: N/A  
Relocation Status: Active  
Relocation Type: Rebalance  
Relocation Rate: Low  
Data to Move Up (GBs): 3.26  
Data to Move Down (GBs): 5.01  
Data to Move Within Tiers (GBs): 2.25  
Data Movement Completed (GBs): 2.00  
Estimated Time to Complete: 4 minutes  
Schedule Duration Remaining: N/A
```

Output for a completed rebalance with an analysis ready for next relocation:

```
Storage Pool Name: Pool 0  
Storage Pool ID: 0  
Relocation Start Time: 02/04/13 11:34  
Relocation Stop Time: 02/04/13 11:34  
Relocation Status: Completed  
Relocation Type: Rebalance  
Relocation Rate: Low  
Data to Move Up (GBs): 7.26  
Data to Move Down (GBs): 5.51  
Data to Move Within Tiers (GBs): 0.00  
Data Movement Completed (GBs): 2.25  
Estimated Time to Complete: 3 minutes  
Schedule Duration Remaining: N/A
```

CHAPTER 5

Compression Commands

This chapter provides a brief overview of the compression feature and includes a list of compression commands that you can use if this feature is enabled on your system.

Major topics include:

- [About LUN Compression](#)342
- [compression -on](#) 342
- [compression -pause](#) 343
- [compression -feature -pause](#) 344
- [compression -resume](#) 344
- [compression -feature -resume](#) 345
- [compression -feature -info](#) 345
- [compression -modify](#) 345
- [compression -off](#) 346
- [compression -list](#) 346

About LUN Compression

Compression is an optional software feature that lets you compress all types of LUN data thereby freeing up storage space in your systems. Data compression analyzes the data on a disk and applies algorithms that reduce the size of repetitive sequences of bits that are inherent in some types of files. The amount of reduction resulting from compression is dependent on the mix of data types being stored; however, unlike de-duplication, compression does not depend on the presence of multiple copies of the same data.

All compression operations perform in the background while continuing to service host I/O. In order to use the compression feature, you must install the compression enabler and the thin provisioning enabler on the system. You can compress classic LUNs as well as pool LUNs. All compressed LUNs are thin LUNs.

During the compression operation for a classic LUN, the software migrates and compresses the LUN data to a thin LUN in a pool. The LUN becomes a compressed thin LUN while migrating the software. During the migration phase, the software detects the zeros on the source LUN and deallocates them on the target LUN which frees up more storage capacity on the target LUN. For better performance and improved space utilization, we strongly recommend that the target LUN is a newly created LUN with no existing data.

Compression operations for pool LUNs (thick and thin) take place within the pool in which the LUN being compressed resides. When compressed, thick and thin LUNs become compressed thin LUNs.

If, for some reason, you need to decompress the compressed data or temporarily pause any active compression operations, the software lets you do both. When decompressed, classic LUNs become fully-provisioned thin LUNs with the same capacity as the original classic LUN and thick LUNs become fully-provisioned thin LUNs. In addition, the software lets you easily monitor the status of all LUN compression operations.

Table 23 naviseccli compression commands

Command	Description
<code>compression -on</code>	Turns on compression on a LUN.
<code>compression -pause</code>	Pauses compression on a LUN.
<code>compression -pauseall</code>	Pauses compression on the system.
<code>compression -resume</code>	Resumes compression on a LUN.
<code>compression -resumeall</code>	Resumes compression on the system.
<code>compression -modify</code>	Modifies the compression rate.
<code>compression -off</code>	Turns off compression.
<code>compression -list</code>	Lists compression information.

compression -on

Turns on compression on a LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -on` command turns on compression for pool LUNs and classic LUNs. The command syntax differs depending on the LUN, either a pool LUN or classic LUN.

SYNTAX

For pool LUNs:

```
naviseccli -h sp compression -on
-l LUN number [-rate low|medium|high|value]
```

For classic LUNs:

```
naviseccli -h sp compression -on
-l LUN number -destpoolid storagepoolid |
-destpoolname storagepoolname
[-rate low|medium|high|value]
```

OPTIONS

`-l LUN number`

Specifies the LUN number to be compressed.

`-destpoolid storagepoolid`

Specifies the destination pool ID to which the classic LUN gets migrated. This option is valid only when the LUN being compressed is a classic LUN.

`-destpoolname storagepoolname`

Specifies the destination pool name to which the classic LUN gets migrated. This option is valid only when the LUN being compressed is a classic LUN.

`-rate low|medium|high|value`

Specifies the compression rate. The default compression rate is medium. The valid compression rates are low or 1, medium or 2, and high or 3.

EXAMPLE

```
naviseccli compression -on -l 10 -destpoolid 0 -rate low
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -pause

Pauses compression on a LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -pause` command pauses compression on a LUN.

SYNTAX

```
naviseccli -h sp compression -pause -l LUN number
```

OPTIONS

`-l LUN number`

Specifies the LUN number of the compressed LUN to be paused.

EXAMPLE

```
naviseccli compression -pause -l 10
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -feature -pause

Pauses compression on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -feature -pause` command pauses compression on the system.

SYNTAX

```
naviseccli -h sp compression -feature -pause [-o]
```

OPTIONS

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli compression -feature -pause
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -resume

Resumes compression on a LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -resume` command resumes paused compression on a LUN. The command can also work for overriding compression in decompression system paused state.

SYNTAX

```
compression -resume -l LUN number [-o]
```

OPTIONS

-l *LUN number*

Specifies the LUN number of the paused compressed LUN to resume.

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli compression -resume -l 10
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -feature -resume

Resumes compression on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -feature -resume` command resumes compression on the system.

SYNTAX

```
naviseccli -h sp compression -feature -resume [-o]
```

OPTIONS

`[-o]`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli compression -feature -resume
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -feature -info

Displays the compression feature information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -feature -info` command displays the compression feature information.

SYNTAX

```
naviseccli -h sp compression -feature -info
```

EXAMPLE

```
naviseccli compression -feature -info
```

OUTPUT

```
Compression Feature State: Paused
```

compression -modify

Modifies the compression rate.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -modify` command changes the compression rate on a specified LUN.

SYNTAX

```
naviseccli -h sp compression -modify -l LUN number -rate low|medium|high|value
```

OPTIONS

-l LUN number

Specifies the LUN number of the compressed LUN to be modified.

-rate low|medium|high|value

Specifies the compression rate. The valid compression rates are low or 1, medium or 2, and high or 3.

EXAMPLE

```
naviseccli compression -modify -l 10 -rate low
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -off

Turns off compression.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -off` command turns off compression on a LUN.

SYNTAX

```
naviseccli -h sp compression -off -l LUN number [-o]
```

OPTIONS

-l LUN number

Specifies the LUN number of the LUN on which you want to stop compression.

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli compression -off -l 10
```

OUTPUT

None if the command succeeds; status or error information if it fails.

compression -list

Lists compression information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli compression -list` command lists compression information on a specified LUN or all LUNs.

SYNTAX

```
naviseccli -h sp compression -list [-l LUN number]
[-destpoolid]
[-uid]
```

```
[-state]
[-status]
[-percentcomplete]
[-rate]
[-totalcapacity]
[-consumedcapacity]
```

OPTIONS

-l *LUN number*

Lists the information about the specified compressed LUN number. If you do not specify the LUN number, all compressed LUNs will be listed.

-destpoolid

Specifies the destination pool ID of the compressed LUN in migrating state. This option is valid only when the compression is in a migrating state.

-uid

Specifies the UID of the compressed LUN.

-state

Specifies the compression state of the LUN.

-status

Specifies the compression status.

-percentcomplete

Specifies the percentage completed.

-rate

Specifies the rate of compression.

-totalcapacity

Specifies the total capacity of the LUN.

-consumedcapacity

Specifies the consumed capacity of the LUN.

EXAMPLE

```
naviseccli compression -list
```

OUTPUT

```
LOGICAL UNIT NUMBER 31
Name: Thin LUN 31
UID: 60:06:01:60:5A:E0:21:00:BA:45:01:A3:68:4F:DE:11
Destination Pool ID: 3
Current State: Migrating
Status: OK(0x0)
Rate: High
Percent Complete: 25
Total Capacity (Blocks): 1234567
Total Capacity (GBs): 10.000
Consumed Capacity (Blocks): 1234567
Consumed Capacity (GBs): 3.500
```


CHAPTER 6

FAST Cache Commands

This chapter provides a brief overview of the EMC FAST Cache feature and includes a list of FAST Cache commands that you can use, if the feature is enabled on your system.

Major topics include:

- [Introduction to caching](#)..... 350
- [cache -fast -create](#)..... 351
- [cache -fast -destroy](#)..... 352
- [cache -fast -info](#)..... 352

Introduction to caching

Caching temporarily stores frequently accessed data in the storage processor's DRAM memory (SP cache) or in FAST Cache. SP cache operates at the finest granularity 8K physical page size with a 64K logical page size. SP cache has the fastest access times, but is limited in size. FAST Cache provides a large secondary cache area with a 64K page size that can be configured in much larger capacities. FAST Cache uses its understanding of LUN mappings to avoid caching data that is already resident on solid-state disk (SSD) storage.

SP cache

SP cache includes read caching and write caching. SP cache dynamically adjusts the amount of cache used by each LUN based on the I/O load on that LUN. SP cache does not allow one LUN to monopolize the whole cache.

- Read caching is always enabled for a system, and the setting cannot be changed.
- Write caching is enabled by default for a system. You can disable write cache at the system level for both classic LUNs and pool LUNs. You can disable write cache at the LUN level for classic LUNs only.

FAST Cache

Note

Storage systems with unbound Flash (SSD) disks and with the FAST Cache enabler installed support the FAST Cache.

FAST Cache is a private storage pool consisting of two or more Flash disks. FAST Cache supports read and write caching. FAST Cache can be enabled for all pool LUNs in a storage pool or for individual classic LUNs.

If the LUN is in a RAID group, it can use FAST Cache if it is enabled for the LUN. If this LUN is in a FAST Cache-enabled pool, it is eligible to use the FAST Cache.

Data already on a Flash drive is not eligible to be moved to FAST Cache.

To use FAST Cache, the system must have:

- FAST Cache enabler installed
- FAST Cache configured
- FAST Cache enabled at the classic LUN and/or storage pool

You can enable FAST Cache for a classic LUN when you create the LUN or after the LUN is bound using the **LUN Cache** or **SP Cache** properties tab. You enable the FAST Cache for a pool LUN when you create the pool or after the pool is created using the **Advanced** pool properties tab.

Note

You should not enable the FAST Cache for write intent log LUNs and Clone Private LUNs. Enabling the FAST Cache for these LUNs is a suboptimal use of the FAST Cache and may reduce the performance of the cache for other LUNs.

naviseccli FAST Cache commands**Table 24** naviseccli FAST Cache commands

Command	Description
<code>cache -fast -create</code>	Creates FAST Cache.
<code>cache -fast -destroy</code>	Destroys FAST Cache.
<code>cache -fast -info</code>	Displays FAST Cache information.

cache -fast -create

Creates FAST Cache.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cache -fast -create` command creates the FAST Cache.

SYNTAX

```
cache -fast -create -disks disksList [-rtype raidtype]
[-mode rw | ro]
[-o]
```

OPTIONS

`-disks disksList`

Specifies the name of the disks used as FAST Cache. If you do not specify the disks, the command will fail.

`-mode rw|ro`

Specifies the cache mode. The `rw` is readwrite mode. The `ro` is readonly mode.

`-rtype raidtype`

Specifies the RAID type. Valid RAID type is `r_1`.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli cache -fast -create -disks 0_0_6 0_0_7 -mode rw -rtype r_1
```

OUTPUT

```
Configuring FAST Cache may take up to several hours to complete. It
will also reduce the amount of Storage System memory available for
use by SP Cache. Do you want to continue? (y/n) y
```

If the user wants to create fast cache and the operation would violate the hot spare policy, then the following confirmation message will be displayed:

```
The selection of disks for the FAST Cache creation will violate the
hot spare policy.
Configuring FAST Cache may take up to several hours to complete. It
will also reduce the amount of Storage System memory available for
use by SP Cache.
Do you want to continue? (y/n) y
```

cache -fast -destroy

Destroys FAST Cache.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cache -fast -destroy` command destroys the FAST Cache.

SYNTAX

```
cache -fast -destroy [-o]
```

OPTIONS

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli cache -fast -destroy
```

OUTPUT

```
To destroy FAST Cache, the System must flush all data to disk. This
operation may be time consuming and may impact system performance.
You can monitor the progress of this operation by using the following
command: cache -fast -info.
Do you want to continue? (y/n) Y
```

cache -fast -info

Displays FAST Cache properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli cache -fast -info` command lets you view the FAST Cache properties.

SYNTAX

```
cache -fast -info [-disks]
[-status]
[-perfData]
```

OPTIONS

`-disks`

Displays the disks used for FAST Cache.

`-status`

Displays FAST Cache status.

`-perfData`

Displays FAST Cache statistics.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE #1**naviseccli cache -fast -info**

```

Disks:
Bus 0 Enclosure 0 Disk 6
Bus 0 Enclosure 0 Disk 7
Mode: Read/Write
Raid Type: r_1
Size (GB): 64
State: Enabled
Current Operation: Initializing
Current Operation Status: Running
Current Operation Percent Completed: 20%
Percentage Dirty SPA: 1
MBs Flushed SPA: 50
Percentage Dirty SPB: 2
MBs Flushed SPB: 80

```

EXAMPLE #2**naviseccli cache -fast -info -status**

```

Mode: Read/Write
Raid Type: r_1
Size (GB): 64
State: Enabled
Current Operation: Initializing
Current Operation Status: Running
Current Operation Percent Completed: 20%

```

EXAMPLE #3**naviseccli cache -fast -info -disks**

```

Disks:
Bus 0 Enclosure 0 Disk 6
Bus 0 Enclosure 0 Disk 7

```

EXAMPLE #4**naviseccli cache -fast -info -perfData**

```

Percentage Dirty SPA: 1
MBs Flushed SPA: 50
Percentage Dirty SPB: 2
MBs Flushed SPB: 80

```

OUTPUT

```

naviseccli cache -fast -info
Disks:
Bus 0 Enclosure 0 Disk 6
Bus 0 Enclosure 0 Disk 7
Mode: Read/Write
Raid Type: r_1
Size (GB): 64
State: Enabled
Current Operation: Initializing
Current Operation Status: Running
Current Operation Percent Completed: 20%
Percentage Dirty SPA: 1
MBs Flushed SPA: 50
Percentage Dirty SPB: 2
MBs Flushed SPB: 80

```


CHAPTER 7

EMC Unisphere QoS Manager

This chapter introduces the EMC® Unisphere™ Quality of Service (QoS) Manager application, associated dependencies, and the online help system.

Major topics are:

- [About Unisphere Quality of Service Manager](#)..... 356
- [QoS Manager CLI commands](#)..... 356
- [Client tools](#)..... 375

About Unisphere Quality of Service Manager

QoS Manager lets you allocate system performance resources on an application-by-application basis. You can use QoS Manager to solve performance conflicts in environments where multiple applications share the same system. Within system capacity, QoS Manager lets you meet specific performance targets for applications, and create performance thresholds to prevent applications from monopolizing system performance.

Note

QoS Manager lets you centrally manage system resources with no host footprint (no host management or software installs are required on the host).

With QoS Manager, you create user-defined I/O classes that characterize, for example, a particular application in order to monitor and control the performance of the I/O class. You must specify various I/O characteristics of your applications to define your I/O classes. You can then determine goals for your I/O classes, and define how to achieve those goals with control methods. You create one or more policies by selecting I/O classes, goals and control methods.

Note

QoS Manager is system specific. Therefore, you must enable QoS Manager on each system for which you want to use QoS Manager. You do not need to reboot your system to enable QoS Manager.

QoS Manager prerequisites

QoS Manager requires the following:

- QoS Manager must be enabled on each system on which you want to use it.

Note

Your system must be running the required EMC VNX OE for Block revision to run QoS Manager. QoS Manager is automatically installed with the required VNX OE for Block revision. You must enable QoS Manager on the system. For required VNX OE for Block revisions, see the release notes supplied with the software.

You do not need to reboot your system to enable QoS Manager.

-
- The correct version of the Java Runtime Environment (JRE) (see the release notes) must be installed on the client from which you run the browser.

For system support and the latest requirements, see the release notes supplied with the software.

QoS Manager CLI commands

This section explains the `naviseccli` command and the CLI commands for QoS Manager application.

Note

The commands in this section function only with a system that has the optional QoS Manager software installed.

nqm -getlatestperf

Displays the latest performance information for a running or measuring policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -getlatestperf` command lets you view the latest performance information for a running or measuring policy.

SYNTAX

```
nqm -getlatestperf [-ioclasses classnames] [-value bw|tt|rt]
```

OPTIONS

`-ioclasses classnames`

Specifies an I/O class name. If not specified, all of the I/O classes in the currently running policy will be listed. If the I/O class specified is not a part of a currently running policy, an error message appears.

`-value bw|tt|rt`

Lists the performance counters output. The following are the valid counter types:

bw: Bandwidth (MB/s)

tt: Throughput (I/O/s)

rt: Response time (ms)

EXAMPLE

For `ss1_spa`, this command displays the latest performance information.

```
naviseccli -h ss1_spa nqm -getlatestperf
```

OUTPUT

```
Class Name:OLTP class
Local Time:Sun June 19 09:10:00 2005
Response Time[ms]:20
Throughput[I/O/sec]:1000
Bandwidth[MB/s]:500
Class Name:Background class
Local Time: Sun June 19 09:10:00 2005
Response Time[ms]:28
Throughput[I/O/sec]:800
Bandwidth[MB/s]:300
```

nqm -info

Provides information about QoS Manager on the system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -info` command provides you with information about QoS Manager on the system.

SYNTAX

```
nqm -info
```

EXAMPLE

For `ssl_spa`, this command displays information about QoS Manager on the system.

```
naviseccli -h ssl_spa nqm -info
```

OUTPUT

```
Current State of NQM: Running
Goal Satisfied: N/A
NQM Provider Version: 6.24
Total User Defined Classes: 2
Total User Defined Policies: 2
Log Event On: Goal no longer satisfied
Fallback Policy:aPolicy
Current Policy: aRunningPolicy
Current Localized Array Time: 22:12:33 10/2/2005
```

nqm -ioclass -create

Creates a new I/O class.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -ioclass` command with the `-create` switch creates a new I/O class. The maximum number of user-defined classes per system is 128.

SYNTAX

```
nqm -ioclass -create -ctrlmethod limit -gmetric bw|tt|rt
-gval value
|-ctrlmethod cruise -gtol value
|-ctrlmethod fixed | -noctrl -name classname
[-snapshots name|WWN | -luns number|WWN | -nolun]
-iotype r|w|rw
[[-minsize value_in_kb][-maxsize value_in_kb]] | [-anyio][-o]
```

OPTIONS

`-ctrlmethod limit`

Limits maximum throughput, maximum bandwidth, or minimum response time.

`-gmetric bw|tt|rt`

Specifies the metric type. It is not valid if the goal type is fixed. Options are:

`bw`: Bandwidth (MB/s)

`tt`: Throughput (I/O/s)

`rt`: Response time (ms)

`-gval value`

Specifies the goal value for this class. If the control method is fixed, this value must be between 1 and 512. Otherwise, it must be greater than 0 and less than 10,000,000,000.

`-ctrlmethod cruise`

Gives the specific value of throughput, bandwidth, or response time.

`-gtol value`

Specifies the tolerance ratio for goal settings. It must be a value between 0 and 100. Unit is percent (%). It is used only for a cruise goal type.

-ctrlmethod *fixed*
 Sets a fixed queue depth for an I/O class.

-noctrl
 If specified, creates the I/O class without any control method.

-name *classname*
 Specifies a new class name.

-snapshots *name/WWN*
 Lists the COFW snapshot LUN names, the snapshot mount point names, or the WWNs.

-luns *number/WWN*
 Lists the LUN numbers or the LUN WWNs. You can use this option for classic LUNs, metaLUNs, thin LUNs, and fully provisioned LUNs.

-nolun
 Specifies that the I/O class will not be modified to any specific LUN.

-iotype *r|w|rw*
 Lists the I/O type. Valid inputs are r (read) or w (write) or rw (readwrite).

-minsize *value_in_kb*
 Indicates the minimum I/O size in KBs, a value between .5 and 64 in increments of .5.

-maxsize *value_in_kb*
 Indicates the maximum I/O size in KBs.

-anyio
 Indicates the I/O class range, if any. If no I/O size range is specified, the default setting will be anyio.

-o
 Creates an I/O class without prompting for warning/confirmation. If you do not specify the -o switch, the following confirmation message appears:

```
You chose Fixed Queue Depth as your control method. Please
be aware that this feature is for EXPERT users ONLY.
Running NQM with this control method may prevent you from
fully utilizing your system's performance resources by
restricting your I/O classes to a certain queue depth on
the array. Please use with caution. Do you want to continue
it (y/n)?
```

EXAMPLE

For `ss1_spa`, this command creates an I/O class with the name `testClass` including LUNs 1, 2, 3 with an I/O type of read. The I/O range is 32–64 KB. The throughput has a goal value of 500 I/Os per second and a tolerance ratio of 10%.

```
naviseccli -h ss1_spa nqm -ioclass -create -name testClass -luns 1 2 3
-iotype r -minsize 32 -maxsize 64 -ctrlmethod cruise -gmetric tt
-gval 500 -gtol 10
```

OUTPUT

None

nqm -ioclass -destroy

Removes an existing I/O class.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -ioclass` command with the `-destroy` switch removes an existing I/O class.

Note

The operation is not allowed while the class is in a running state.

SYNTAX

```
nqm -ioclass -destroy -name classname [-o]
```

OPTIONS

`-name classname`

Specifies an existing class name.

`-o`

Destroys an I/O class without prompting for warning/confirmation. If you do not specify the `-o` switch when destroying an I/O class, a confirmation message appears.

y = Destroy an I/O class.

n = Do not destroy an I/O class.

EXAMPLE

For `ssl_spa`, `testClass` will be removed.

```
naviseccli -h ssl_spa nqm -ioclass -destroy -name testClass
```

OUTPUT

None

nqm -ioclass -list

Lists all application classes previously created.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -ioclass` command with the `-list` switch displays all previously created application classes.

SYNTAX

```
nqm -ioclass -list [-ctrlmethod]
[-iorange]
[-iotype]
[-luns]
[-name [ioclassname]]
[-snapshots]
[-status]
[-state]
```

OPTIONS

`-ctrlmethod`

Displays the control method and one or more of the following parameters. For the cruise control method, it displays the goal metric, goal value, and tolerance. For the limit method, it displays the goal metric and goal value. For the fixed-queue-depth method, it displays only the queue depth value.

`-iorange`

Lists the range of I/O sizes in this class.

`-iotype`

- Lists the I/O type. They are read, write, or readwrite.
- luns
 - Lists the LUNs contained in the class. For each classic LUN, thin LUN, and fully provisioned LUN, the LUN number, name, WWN, and RAID type are displayed. For metaLUNs, the RAID type is displayed as N/A. For the class that does not contain any LUNs, the LUNs are listed as Not LUN Specific.
- name [*ioclassname*]
 - Specifies an I/O class name and lists all information for that I/O class. If you do not specify this option, all I/O class names are displayed.
- snapshots
 - Lists the snapshots LUN (s) and the snapshot mount points contained in the class. For each snapshot LUN, the snapshot LUN name, the WWN, and the target logical unit are displayed.
- status
 - Displays the current status of an I/O class, which may be one of the following:
 - Ready: The I/O class is ready to be run within the context of a policy.
 - Warning: The I/O class does not have a goal and can only be measured in the context of a policy.
 - Broken: The I/O class does not include any LUNs because its last LUN, which could be a classic LUN, a metaLUN, or a thin LUN, was destroyed. If the last LUN is a thin LUN and it goes offline, the I/O class (or any policies that contain this I/O class) is displayed as Broken.
- state
 - Displays the class operational state, which may be one of the following:
 - Running: QoS Manager actively controls the system's I/O flow and collects performance information.
 - Idle: QoS Manager is not controlling or measuring I/O.
 - Measuring: QoS Manager collects performance, but does not tune it.

EXAMPLE

For `ss1_spa`, this command displays all previously created application classes.

```
naviseccli -h ss1_spa nqm -ioclass -list
```

OUTPUT

```
Name: TestClassWithNOGoal
Status: Ready
Current State: Idle
Number of LUNs: 3
LUNs:
LUN Number: 0
LUN Name: LUN 0
LUN WWN: 60:06:01:60:4E:22:0E:00:4C:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
LUN Number: 2
LUN Name: LUN 2
LUN WWN: 60:06:01:60:4E:22:0E:00:4E:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
LUN Number: 1
LUN Name: LUN 1
LUN WWN: 60:06:01:60:4E:22:0E:00:4D:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
Number of Snapshot(s): 1
Snapshot(s):
Snapshot Name: Snap7
Snapshot WWN: 60:06:01:60:4B:50:1D:00:B4:21:88:CF:39:E6:DD:11
```

```

Target Logical Unit: 8170
IO Type: ReadWrite
IO Size Range: Any
Control Method: Limit
Goal Metric: Throughput
Metric Type: Throughput
Goal Value: 400 IO/s

Name: TestCompleteClass
Status: Ready
Current State: Idle
Number of LUNs: 3
LUNs:
LUN Number: 0
LUN Name: LUN 0
LUN WWN: 60:06:01:60:4E:22:0E:00:4C:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
LUN Number: 2
LUN Name: LUN 2
LUN WWN: 60:06:01:60:4E:22:0E:00:4E:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
LUN Number: 1
LUN Name: LUN 1
LUN WWN: 60:06:01:60:4E:22:0E:00:4D:CC:D6:F8:67:6C:D8:11
RAID Type: RAID5
IO Type: ReadWrite
IO Size Range: 12KB -- 32KB
Control Method: Cruise Control
Goal Metric: Throughput
Metric Type: Throughput
Goal Value: 400 IO/s
Goal Tolerance[%]: 50

```

nqm -ioclass -modify

Modifies an existing I/O class.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -ioclass` command with the `-modify` switch modifies an existing I/O class. You can use this command to rename the class, change the properties of a class, or update the contained target lists.

Note

This operation is not allowed while the class is in a running state.

SYNTAX

```

nqm -ioclass -modify -name classname
[-anyio]
[-iotype r|w|rw]
[-snapshots name|WWN|-luns number|WWN|-nolun]
[[-maxsize value_in_kb] [-minsize value_in_kb]]
[-newname newname]
[goal switches]
[-o]
[-ctrlmethod limit -gmetric bw|tt|rt -gval value|
-ctrlmethod cruise
-gmetric bw|tt|rt -gval value -gtol value|
-ctrlmethod fixed -gval value|-noctrl]

```

OPTIONS

`-name classname`

Specifies an existing class name.

Note

`-anyio` and `-maxsize/-minsize` are exclusive options.

`-anyio`

Indicates the range of any I/O class.

`-iotype r|w|rw`

Lists the I/O type. Valid inputs are r (read) or w (write) or rw (readwrite).

`-snapshots name/WWN`

Lists the COFW snapshot LUN names, the snapshot mount point names, or the WWNs.

`-luns number/WWN`

Lists LUN numbers or LUN WWNs. You can use this option for classic LUNs, metaLUNs, thin LUNs, and fully provisioned LUNs.

`-nolun`

Specifies that the I/O class will not be modified to any specific LUN.

`-maxsize value_in_kb`

Indicates the maximum I/O size in KBs.

`-minsize value_in_kb`

Indicates the minimum I/O size in KBs, a value between .5 and 64 in increments of .5.

`-newname newname`

Indicates a new I/O class name.

`-o`

Modifies an I/O class without prompting for warning or confirmation. If the user omits this option when modifying an I/O class, a confirmation message appears:

```
The I/O class is contained by a scheduled policy. Do you
want to modify it (y/n)?
```

`-ctrlmethod limit`

Limits maximum throughput, maximum bandwidth, or minimum response time.

`-gmetric bw|tt|rt`

Specifies the metric type. It is not valid if the goal type is fixed. The following are the valid metric types:

`bw`: Bandwidth (MB/s)

`tt`: Throughput (I/O/s)

`rt`: Response time (ms)

`-gval value`

Specifies the goal value for this class. If the control method is fixed, this value must be between 1 and 512. Otherwise, it must be greater than 0 and less than 10,000,000,000.

`-ctrlmethod cruise`

Gives the specific value of throughput, bandwidth, or response time.

`-gtol value`

Specifies the tolerance ratio for goal settings. It must be a value between 0 and 100. The unit is percent (%). It is used only for a cruise goal type.

`-ctrlmethod fixed`

Sets a fixed queue depth for an I/O class.

`-noctrl`

If specified, creates the I/O class without any control method.

EXAMPLE

For `ss1_spa`, this command modifies `testClass` and gives a new I/O class name of `OLTPClass`.

```
naviseccli -h ss1_spa nqm -ioclass -modify -name testClass -newname
OLTPClass
```

OUTPUT

None

nqm -measure

Starts a specified policy in measuring mode.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -measure` command starts a specified policy in measuring mode.

SYNTAX

```
nqm -measure policyname
```

OPTIONS

policyname

Specifies the existing policy name.

EXAMPLE

For `ss1_spa`, this command starts `testPolicy` in measuring mode.

```
naviseccli -h ss1_spa nqm -measure testPolicy
```

OUTPUT

None

nqm -off

Shuts down the QoS Manager operations.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `nqm -off` command shuts down the QoS Manager operations. Use this command to stop QoS Manager from running, suspend all scheduled tasks, and disable the fallback feature. This command does not work if QoS Manager is not active.

SYNTAX

```
nqm -off [-o]
```

OPTIONS

-o

Shuts down the QoS Manager operations without prompting for warning/confirmation. If you omit the `-o` switch, the following confirmation message appears:

```
This command will stop NQM from running, suspend all
scheduled tasks and disable fallback feature. Do you want
to continue (y/n)?
```

EXAMPLE

For 10.15.17.205, this command shuts down QoS Manager operations.

```
naviseccli -user admin -password admin -scope 0 -h 10.15.17.205 nqm
-off
```

OUTPUT

None

nqm -policy -create

Creates a new policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -policy` command with the `-create` switch creates a new policy. A maximum of 10 policies are supported on a storage system.

SYNTAX

```
nqm -policy -create -name policyname [-evalwindow
number_of_samples]
[-failaction stop|measure|continue]
[-ioclasses classnames|-noclass]
[-timelimit value_in_min]
```

OPTIONS

`-name policyname`

Specifies the policy name.

`-evalwindow number_of_samples`

Specifies the number of sample periods needed to meet (or to fail to meet) the goal in order for the goal to be classified as satisfied (or unsatisfied). This value must be between 3 and 30. The default value is 3.

`-failaction stop|measure|continue`

Specifies the action QoS Manager takes when the goal is not satisfied. The default action is stop.

`-ioclasses classnames`

Specifies the I/O class names. The maximum I/O classes allowed is 64.

`-noclass`

Creates the policy with no I/O class.

`-timelimit value_in_min`

Specifies the maximum time in minutes that a system takes to satisfy the goal. You must specify the limit in minutes. This value must be between 5 and 1440. The default value is 60.

EXAMPLE

For `ss1_spa`, this command creates a new policy for `clitestpolicy`. The specified I/O class is `testClass`. The system has 20 minutes to satisfy the goal. QoS Manager stops if the goal is not satisfied in the given time limit. The number of sample periods is 3.

```
naviseccli -h ss1_spa nqm -policy -create -name clitestpolicy
-ioclasses testClass -timelimit 20 -failaction stop -evalwindow 3
```

OUTPUT

None

nqm -policy -destroy

Deletes an existing policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -policy` command with the `-destroy` switch deletes an existing policy that is not contained by any scheduled task. This command will not be allowed if the specified policy is running or measuring.

SYNTAX

```
nqm -policy -destroy -name policyname [-o]
```

OPTIONS

`-name policyname`

Specifies the policy name to be destroyed.

`-o`

Destroys the policy without prompting for warning/confirmation. If you do not specify the `-o` switch when destroying a policy, a confirmation message appears.

`y` = Destroy a policy.

`n` = Do not destroy a policy.

EXAMPLE

For `ss1_spa`, this command deletes the policy, `testPolicy`.

```
naviseccli -h ss1_spa nqm -policy -destroy -name testPolicy
```

OUTPUT

None

nqm -policy -list

Creates and manages QoS Manager policies.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -policy` command with the `-list` switch creates and manages the QoS Manager policies.

SYNTAX

```
nqm -policy -list [-evalwindow]
[-failaction]
[-ioclasses classnameses]
[-name [policyname]]
[-state]
[-status]
[-timelimit]
```

OPTIONS

`-evalwindow`

- Specifies the number of sample periods needed to meet (or to fail to meet) the goal in order for the goal to be classified as satisfied (or unsatisfied).
- failaction
Specifies the action QoS Manager takes when it fails to satisfy the goal. If the policy contains a class with the control method of fixed queue depth, this switch option will be N/A.
- ioclasses *classnames*
Specifies the I/O class names.
- name [*policyname*]
Specifies the policy name. If not specified, all policy names will be displayed.
- state
Displays the policy's current state.
- status
Displays the policy's current status.
- timelimit
Specifies the maximum time it should take the system to satisfy the goal. The unit is measured in minutes.

EXAMPLE

For `ssl_spa`, the command displays the created QoS Manager policies.

```
naviseccli -h ssl_spa nqm -policy -list
```

OUTPUT

```
Name: TestInCompletePolicy2
Status:Ready
Current State: Idle
Number of I/O Classes: 2
I/O Classes:
Name: TestClassWithNOGoal
Name: Background Class
Action on Failure: Stop NQM tuning after failure to converge
Time Limit:15 min
Evaluation Window: 10 samples

Name: TestCompletePolicy
Status:Ready
Current State: Idle
Number of I/O Classes: 2
I/O Classes:
Name: TestCompleteClass
Name: Background Class
Action on Failure: Stop NQM after failure to satisfy goal
Time Limit:15min
Evaluation Window: 10 samples
```

nqm -policy -modify

Modifies an existing policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -policy` command with the `-modify` switch modifies an existing policy.

SYNTAX

```
nqm -policy -modify -name policyname [-evalwindow
number_of_samples]
[-failaction stop|measure|continue]
```

```
[-ioclasses classnames] | [-noclass]
[-newname newname]
[-o]
[-timelimit value_in_min]
```

OPTIONS

-name *policyname*
Specifies the policy name.

-evalwindow *number_of_samples*
Specifies the number of sample periods needed to meet (or to fail to meet) the goal in order for the goal to be classified as satisfied (or unsatisfied). This value must be between 3 and 30. The default value is 3.

-failaction stop |measure|continue
Specifies the action QoS Manager takes when the goal is not satisfied. The default action is stop.

-ioclasses *classnames*
Specifies the I/O class names.

-noclass
Modifies the policy to contain no I/O classes.

-newname *newname*
New name for an existing policy.

-o
Modifies a policy without prompting for warning/confirmation.

-timelimit *value_in_min*
Specifies the maximum time it should take the system to satisfy the goal. Measured in minutes, the value must be between 5 and 1440. The default value is 60.

EXAMPLE

For `ss1_spa`, this command modifies the existing policy name, `clitestpolicy2`, and assigns a new name, `OLTPPolicy`. The QoS Manager stops if the goal is not satisfied.

```
naviseccli -h ss1_spa nqm -policy -modify -name clitestpolicy2
-newname OLTPPolicy -failaction stop
```

OUTPUT

None

nqm -run

Starts a specified policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -run` command starts a specified policy.

SYNTAX

```
nqm -run policyname
```

OPTIONS

policyname
Specifies which policy to run.

EXAMPLE

For `ss1_spa`, this command runs `testPolicy`.

```
naviseccli -h ss1_spa nqm -run testPolicy
```


OUTPUT

None

nqm -schedule -create

Creates a scheduled task.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-create` switch creates a scheduled task.

SYNTAX

```
nqm -schedule -create -name schedulename -startdate mm/dd/yyyy
-starttime
hh:mm:ss -duration hours:minutes -policy polycyname
[-daily frequency]
| [-weekly frequency -pattern [su] [mo] [tu] [we] [th] [fr]
[sa]] | -norecur
```

OPTIONS

- name *schedulename*
Specifies the scheduled task name.
- startdate *mm/dd/yyyy*
Indicates the date the task is scheduled to start.
- starttime *hh:mm:ss*
Indicates the time the task is scheduled to start.
- duration *hours:minutes*
Indicates how long the policy will run. The hours can be a value between 0 and 168.
Minutes is a value between 0 and 59.
- policy *polycyname*
Specifies the policy name to which the scheduled task is associated.
- daily *frequency*
Specifies that this scheduled task should occur daily at the specified frequency.
- weekly *frequency*
Specifies that this scheduled task should occur weekly at the specified frequency.
- pattern [su] [mo] [tu] [we] [th] [fr] [sa]
Specifies the list of weekdays, separated by spaces, that the task will reoccur.
Specify this switch only with the `-weekly` switch.
- norecur
Specifies no recurrence of the scheduled task, if specified.

EXAMPLE

For `ssl_spa`, this command creates a daily recurring task for `TestScheduleRecurring`. The task will occur every day at 12:10:00. The duration will be 4 hours and 20 minutes, effective 4/5/2004. The scheduled task will be associated with `TestCompletePolicy`.

```
naviseccli -h ssl_spa nqm -schedule -create -name
TestScheduleRecurring -startdate 04/05/2004 -starttime
12:10:00 -duration 4:20 -policy TestCompletePolicy -daily 1
```

OUTPUT

None

nqm -schedule -destroy

Clears a specified scheduled task.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-destroy` switch clears a specified scheduled task.

SYNTAX

```
nqm -schedule -destroy -name schedulename [-o]
```

OPTIONS

`-name schedulename`

Specifies the scheduled task name.

`-o`

Destroys the scheduled task without prompting for warning/confirmation. If you do not specify the `-o` switch when destroying a scheduled task, a confirmation message appears.

y = Destroy a scheduled task.

n = Do not destroy the scheduled task.

EXAMPLE

For `ss1_spa`, this command clears the scheduled task `testSchedule`.

```
naviseccli -h ss1_spa nqm -schedule -destroy -name testSchedule
```

OUTPUT

None

nqm -schedule -list

Lists all scheduled tasks.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-list` switch lists all scheduled tasks in a timely order. If there is no task scheduled, then no task will be displayed.

SYNTAX

```
nqm -schedule -list [-name [schedulename]]
[-state]
[-start]
[-duration]
[-policy policyname]
[-recur]
```

OPTIONS

`-name schedulename`

Lists all information for a specific QoS Manager task.

`-state`

Specifies the state of the tasks. The different states include:

scheduled: A scheduled occurrence is waiting for its turn to run.

suspended: A suspended task is temporarily excluded and will not be run at its scheduled start time.

running: A scheduled policy is currently running.

broken: A broken task contains a broken policy that cannot be run.

-start

Displays the calendar time to start QoS Manager.

-duration

Specifies the duration the schedule runs.

-policy *polycyname*

Specifies the name of the policy to be run.

-recur

Displays the recurring pattern or no recurrence.

EXAMPLE

For `ssl_spa`, this command lists all scheduled tasks.

```
naviseccli -h ssl_spa nqm -schedule -list
```

OUTPUT

```
Name: TestScheduleRecurring
Current State: Scheduled
Next Start Time: Mon Apr 05 10:10:00 2004
Duration: 2 hours
Policy Name: TestCompletePolicy
Recurrence Type: Weekly
Frequency: Every 2 weeks
Pattern: Sunday, Monday and Tuesday

Name: TestScheduleRecurring
Current State: Running
Next Start Time: Mon Apr 05 12:00:00 2004
Duration: 4 hours
Policy Name: TestCompletePolicy
Recurrence Type: Weekly
Frequency: Every week
Pattern: Sunday

Name: TestScheduleRecurringDaily
Current State: Idle
Next Start Time: Mon Apr 05 16:00:00 2004
Duration: 2 hours
Policy Name: TestCompletePolicy
Recurrence Type: Daily
Frequency: Every day
Pattern: N/A

Name: TestScheduleNoRecurring
Current State: Idle
Next Start Time: Mon Apr 05 18:10:00 2004
Duration: 2 hours
Policy Name: TestCompletePolicy
Recurrence Type: None
Frequency: N/A
Pattern: N/Az
```

nqm -schedule -modify

Modifies a scheduled task.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-modify` switch modifies a specified scheduled task.

SYNTAX

```
nqm -schedule -modify -name schedulename [-newname newname]
[-policy policyname]
[-startdate mm/dd/yyyy ]
[-starttime hh:mm:ss]
[-duration hours:minutes ]
[-daily frequency] | [-weekly frequency -pattern [su] [mo] [tu]
[we] [th] [fr] [sa]] | -norecur
```

OPTIONS

`-name schedulename`

Specifies the existing scheduled task name.

`-newname newname`

Specifies a new name for the existing scheduled task.

`-policy policyname`

Indicates the policy name to which the scheduled task is associated.

`-startdate mm/dd/yyyy`

Indicates the date the task is scheduled to start.

`-starttime hh:mm:ss`

Indicates the time the task is scheduled to start.

`-duration hours:minutes`

Indicates the length of time the policy will run.

`-daily frequency`

Specifies that this scheduled task should occur daily at the specified frequency.

`-weekly frequency`

Specifies that this scheduled task should occur weekly at the specified frequency.

`-pattern [su] [mo] [tu] [we] [th] [fr] [sa]`

Specifies the list of weekdays, separated by spaces, that the task will reoccur. This switch should be specified only with the `-weekly` switch.

`-norecur`

Creates the scheduled task with no recurrence.

EXAMPLE

For `ssl_spa`, this command modifies the task `testSchedule`. The task will start on 6/11/2004 at 12:00:00 and will run for 5 hours and 30 minutes.

```
naviseccli -h ssl_spa nqm -schedule -modify -name testSchedule
-startdate 6/11/2004 -starttime 12:00:00 -duration 5:30
```

OUTPUT

None

nqm -schedule -resume

Resumes a scheduled task.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-resume` switch resumes a scheduled task.

SYNTAX

```
nqm -schedule -resume -all | -name schedulename
```

OPTIONS

`-all`

Resumes all existing tasks.

`-name schedulename`

Specifies an existing scheduled task name.

EXAMPLE #1

For `ss1_spa`, this command resumes the scheduled task `testSchedule`.

```
naviseccli -h ss1_spa nqm -schedule -resume -name testSchedule
```

EXAMPLE #2

For `ss1_spa`, this command resumes all existing scheduled tasks.

```
naviseccli -h ss1_spa nqm -schedule -resume -all
```

OUTPUT

None

nqm -schedule -suspend

Suspends scheduled tasks.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -schedule` command with the `-suspend` switch suspends a scheduled task.

SYNTAX

```
nqm -schedule -suspend -all | -name schedulename [-o]
```

OPTIONS

`-all`

Suspends all existing tasks.

`-name schedulename`

Specifies the existing scheduled task.

`-o`

Suspends a scheduled task without prompting for confirmation.

EXAMPLE #1

For `ss1_spa`, this command suspends the task named `testSchedule`.

```
naviseccli -h ss1_spa nqm -schedule -suspend -name testSchedule
```

EXAMPLE #2

For `ss1_spa`, this command suspends all scheduled tasks without prompting for confirmation.

```
naviseccli -h ss1_spa nqm -schedule -suspend -all -o
```

OUTPUT

None

nqm -setfallback

Sets or unsets a fallback policy.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -setfallback` command will set or unset a fallback policy.

SYNTAX

```
nqm -setfallback -enable pname | -disable
```

OPTIONS

`-enable pname`

Enables and sets a fallback policy with the specified policy name.

`-disable`

Disables and unsets the fallback policy.

EXAMPLE

For `ss1_spa`, this command enables a fallback policy for `apolicy`.

```
naviseccli -h ss1_spa nqm -setfallback -enable apolicy
```

OUTPUT

None

nqm -setoptions

Toggles data logging and event control switches.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -setoptions` command toggles data logging and event control switches.

SYNTAX

```
nqm -setoptions -logevent [[gs][gu][gn]] | none
```

OPTIONS

`-logevent [[gs][gu][gn]] | none`

Indicates if an event needs to be added in the event log due to the following four conditions:

`gs`: Goal satisfied

`gu`: Goal unattainable

`gn`: Goal no longer satisfied

`none`: Log no events

EXAMPLE

For `ssl_spa`, this command toggles the data logging and the event control switch. Events will be put into the event log if the goal is satisfied, unattainable, or no longer satisfied.

```
naviseccli -h ssl_spa nqm -setoptions -logevent gs gu gn
```

OUTPUT

None

nqm -stop

Stops the active policy immediately.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli nqm -stop` command stops the active policy (running or measuring) immediately.

SYNTAX

```
nqm -stop [-nofallback] [-o]
```

OPTIONS

`-nofallback`

Starts the fallback policy unless this switch is specified.

`-o`

Stops a policy without prompting for warning/confirmation.

EXAMPLE

For `ssl_spa`, the command immediately stops the active policy without starting the fallback policy.

```
naviseccli -h ssl_spa nqm -stop -nofallback
```

OUTPUT

None

Client tools

This section describes the QoS Manager client tools.

nqm -archive -dump

Dumps QoS Manager archive files.

PREREQUISITES

The host, username/password, and scope options of `naviseccli` are ignored since this command runs locally.

DESCRIPTION

The `naviseccli nqm` command with the `-archive -dump` command dumps the performance data, relationship, configuration, and statistics information in a format that you specify. You can specify the output format, XML or CSV, for the dumped information.

The `-archive -dump` command includes the following subcommands:

- `-data`

- -rel
- -config
- -stats

Note

The `-archive -dump` command recognizes and reports an error message when you try to open an encrypted archive file.

SYNTAX

```
nqm -archive -dump -data filename(s) [-out filename(s)]
[-join]
[-overwrite y|n]
[-xml|-delim code]
[-enquote code]
[-eol code]
[-header y|n]
[-object codes] [-format codes] [-stime starttime]
[-ftime endtime] [-timezone code] [-progress]
```

or

```
nqm -archive -dump -rel filename(s) [-out filename(s)]
[-join]
[-overwrite y|n]
[-xml|-delim code]
[-enquote code]
[-eol code]
[-root objcode]
[-level depth]
[-progress]
```

or

```
nqm -archive -dump -config filename(s) [-out filename(s)]
[-join]
[-overwrite y|n]
[-xml] [-delim code] [-enquote code]
[-eol code] [-object objcode] [-progress]
```

or

```
nqm -archive -dump -stats filename(s) [-out filename(s)]
[-join] [-overwrite y|n] [-timezone code]
```

OPTIONS

`-data filename(s)`

Dumps data from archive files. You must provide the archive filename following the `-data` command.

Note

It lets you specify QoS Manager archive files and dump QoS Manager specific data.

The optional switches are as follows:

`-out filename(s)`

Dumps multiple files. Specify multiple output filenames separated by commas. The output filenames list corresponds to the dumped filenames list.

Note

You can dump multiple files into a single file using the `-join` switch.

-join

Appends the output from multiple files into a single file. The header will be dumped only once if this switch is used and if the header switch is enabled. This switch is valid only if you specify multiple files.

-overwrite y|n

Works with the `-out` switch. If you use the `-overwrite` switch without the `-out` switch, it is ignored. The `-overwrite` switch applies to all the output files if you specify multiple files to dump.

y = Without a warning, overwrites the existing file specified by the `-out` switch.

n = Displays an error message if the file that the `-out` switch specified already exists.

Note

If you try to overwrite a read-only file, an error message appears.

-xml

Specifies the output as XML format. This switch is optional, and you cannot use it with the `-delim` type. If you do not specify the `-xml` switch, delimited output results. If you use it with the `-data` command and the `-format` switch, an error message appears.

-delim code

Specifies a delimiter that separates the output data. With this switch, you can specify the output as delimited type. It is optional and cannot be used with the `-xml` switch. Specify the delimiter when you use this switch. If you do not specify either `-xml` or `-delim`, delimited output results.

Table 25 Delimiter codes

Delimiter	Code
Comma (default)	cm
Semicolon	sc
Colon	cl
Tab	tb
Slash	sl
Backslash	bs
Period	pd
White space	ws

Note

If you use a slash or a period as a delimiter, you must also use the `-enquote` switch.

-enquote code

Encloses each performance characteristic within a set of special characters. If you specify the output type as `-xml`, an error message appears.

The possible special characters that you can use for the `-enquote` switch are listed below:

Table 26 Valid `-enquote` switch special characters

Enclosing character	Code
None	no
Quote	qu
Apostrophe	ap
Accent character	ac

Note

Use this switch only when you specify an output type as `-delimited`.

`-eol code`

Specifies a particular end-of-line character. The output reaches an end of line when all performance characteristics for one specific object at a particular time have been output. The underlying client system determines the default end-of-line character. If you specify the output type as `-xml`, an error message appears. The possible end-of-line characters for the switch are listed below:

Table 27 Valid end-of-line characters

End-of-line character	Code
Carriage return	cr
Line feed	lf

Note

Use this switch only when you specify an output type as `-delimited`.

`-header y|n`

Specifies whether to output a column header.

y = Output a header

n = Do not output a header

`-object codes`

Specifies the objects for which to collect the performance statistics and the order in which to output them. Not specifying this switch dumps all policy session objects.

Table 28 Valid object type codes

Object type	Code
Policy session (QoS Manager)	ps

`-format codes`

Specifies which performance characteristics to output and the order in which they will appear. This switch is optional and cannot be used with the `-xml` switch. If you use the `-xml` switch with the `-format` switch, an error message appears. The `-format` switch requires one or more of the codes listed below. If you do not specify this switch, all performance characteristics are output in the order listed below:

Table 29 Performance characteristics without `-format` code

Performance characteristic	Object type	Code
Object Name	All objects	on
Owner Array Name	All objects	oan
Poll Time	All objects	pt
Policy Start Time	Policy Session	psst
Policy End Time	Policy Session	pset
I/O Class Name	Policy Session	cn
I/O Delay	I/O Class	fd
Response Time (ms)	LUN, metaLUN, disk, SP, I/O Class, Snapshot, Thin LUN	rt
Total Bandwidth (MB/s)	LUN, metaLUN, disk, SP, I/O Class, Snapshot, Thin LUN	tb
Total Throughput (I/O/s)	LUN, metaLUN, disk, SP, I/O Class, Snapshot, Thin LUN	tt
Queue Depth	LUN, metaLUN, disk, SP, I/O Class, Snapshot, Thin LUN	qd

`-stime starttime`

Specifies an archive poll time to start iterating through the data. The `-stime` switch requires the start time to be in the following format:

mm/dd/yyyy hh:mm:ss

where:

mm is the month of a year (1-12)

dd is the day of a month (1-31)

yyyy is the four-digit year

hh is the hour of the day (0-23)

mm is the minute in an hour (0-59)

ss is the second in a minute (0-59).

You must enclose the time in double quotation marks.

The `-stime` switch works in combination with the `-ftime` switch, described below. If the start time is greater than the end time, an error message appears. If you do not specify a start time, this switch defaults to the earliest poll time in the archive file.

Note

If you use the `-stime` switch with the `-timezone` switch, you must apply the time zone to the start time.

If the start time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the start time format is not valid, an error message appears.

`-ftime endtime`

Specifies an archive poll time to stop iterating through the data. The `-ftime` switch requires the end time to be in the following format:

`mm/dd/yyyy hh:mm:ss`

where:

mm is the month of a year (1-12)

dd is the day of a month (1-31)

yyyy is the four-digit year

hh is the hour of the day (0-23)

mm is the minute in an hour (0-59)

ss is the second in a minute (0-59).

You must enclose the time in double quotation marks.

The `-ftime` switch works in combination with the `-stime` switch, described above. If the start time is greater than the end time, an error message appears. If you do not specify an end time, this switch defaults to the latest poll time in the archive file.

Note

If you use the `-ftime` switch with the `-timezone` switch, you must apply the time zone to the end time.

If the end time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the end time format is not valid, an error message appears.

`-timezone code`

Applies a particular time zone to the performance characteristic, Poll Time, included in the output data. By default, QoS Manager displays this characteristic in the local time of the client system. The possible time zones for the switch are listed below:

Table 30 Time zone values

Time zone	Code
Local time (default)	local
Universal time	utc

`-progress`

Displays information about the current dump status, including the name of the file being dumped and the percentage dumped. An approximate value is calculated based on the number of objects for the percentage dumped information. This switch works with the `-out` switch. If you use the `-progress` switch without the `-out` switch, it is ignored. Do not use this switch in scripts.

`-rel`

Dumps the relationship information for the various objects in the archive file. Unisphere version 1.0 contains fully provisioned LUN support for QoS Manager. You can dump and display the fully provisioned LUN if it is part of any I/O class. You can

also dump and display the snapshot mount point if it is a part of any I/O class. You must provide the archive filename following the `-rel` command. You can specify multiple archive filenames separated by commas to dump the relationship information for all of them.

In addition to the switches, `-join`, `-out`, `-overwrite`, `-progress`, `-eol`, and `-enquote` (which are described above), you can also use the following optional switches with the `-rel` command:

`-xml`

Specifies the output in XML format. This switch is optional and you cannot use it with the `-delim` type. If you do not specify the `-xml` switch, delimited output results.

`-delim`

Specifies the output as delimited. In the delimited format, the relationship is dumped in a logical tree format, whereby each node appears on a separate line.

`-config`

Dumps the configuration information of objects in the archive file. You must provide the archive filename following the `-config` command. You can specify multiple archive files separated by commas to dump the configuration information for all of them. In addition to the switches, `-join`, `-out`, `-overwrite`, `-progress`, `-delim`, `-enquote`, and `-eol` (which are described above), you can also use the following optional switches with the `-config` command:

`-object objcode`

Specifies which objects and in what order you want to dump the configuration information. If you do not specify this switch, all object types are dumped in the order of policies, I/O classes, metaLUNs, classic LUNs, thin LUNs, fully provisioned LUNs, and snapshot mount points. If you use the `plu` object code for pool-based LUNs, the `tl` option for thin LUNs is ignored. Use the `al` object code to dump all LUN targets including metaLUNs, classic LUNs, thin LUNs, fully provisioned LUNs, and snapshot LUNs. All other object codes are ignored if you specify the `al` object code.

Table 31 Order of dump configuration information

Object type	Code
Policy session	ps
I/O class	cl
All LUNs (including all supported public LUNs)	al
MetaLUN	ml
Classic LUN	l
Thin LUN	tl
Pool LUN (including thin LUN and fully provisioned LUN)	plu
Snapshot LUN	sl

`-xml`

Specifies the output in XML format. This switch is optional and you cannot use it with the `-delim` type. If you do not specify the `-xml` switch, delimited output results.

`-stats archive filename(s)`

Dumps the statistical information for multiple archive files. A blank line separates dumped information of different files. You can use the following optional switches

with the `-stats` command: `-join`, `-out`, `-overwrite` (which are described above).

EXAMPLE #1

This command lets you dump data from the archive file `archive.nqm` without outputting a header.

```
naviseccli -h ssl_spa nqm -archive -dump -data archive.nqm -header n
```

EXAMPLE #2

This command lets you dump relationship information of various objects in the archive file `archive.nqm` and output it in XML format.

```
naviseccli -h ssl_spa nqm -archive -dump -rel archive.nqm -xml
```

EXAMPLE #3

This command lets you dump configuration information of objects in the archive file `archive.nqm`; if the output file `archive.csv` exists, the system overwrites it without a warning.

```
naviseccli -h ssl_spa nqm -archive -dump -config archive.nqm -out archive.csv -overwrite y
```

EXAMPLE #4

This command lets you dump statistical information for multiple archive files `archive1.nqm` and `archive2.nqm`.

```
naviseccli -h ssl_spa nqm -archive -dump -stats archive1.nqm, archive2.nqm
```

OUTPUT

Varies depending upon which switch you use.

nqm -archive -retrieve

Retrieves QoS Manager archive files.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Note

The `navicli nqm` command is not required for the `-archive -retrieve` command to run.

DESCRIPTION

The `naviseccli nqm` command with the `-archive -retrieve` command lets you create and retrieve the archive files from the system.

SYNTAX

```
nqm -archive -retrieve -file filename [-location directorypath]
[-overwrite y|n]
[-retry number]
[-v]
```

OPTIONS

`-file filename`

Specifies the name of the archive file once it is retrieved to the client system.

`-location directorypath`

Specifies the path location of the archive file once it is retrieved to the client system. If you do not specify a location, QoS Manager uses the current working directory.

`-overwrite y|n`

Specifies whether to overwrite an existing archive file on the client system. If the specified archive file already exists and this switch is not used, the command fails.

`y` = Without a warning, overwrites the existing file specified by the `-out` switch.

`n` = Displays an error message if the file that the `-out` switch specified already exists.

`-retry number`

Specifies the number of times to retry a retrieval before giving up completely. A retry will occur only when the actual move of the archive file from the system to the client system fails.

`-v`

As the retrieval process progresses, the status of each stage will be output.

EXAMPLE

This command retrieves an archive to the file, `archive.nqm` in `c:\Temp`, overwrites the existing file and retries two times.

```
naviseccli -h ssl_spa nqm -archive -retrieve -file archive.nqm
-location c:\temp -overwrite y -retry 2
```

OUTPUT

None

CHAPTER 8

EMC Unisphere Analyzer

This chapter introduces the EMC® Unisphere™ Analyzer performance monitor application, which measures the performance of disk-array storage systems.

Major topics are:

- [About Unisphere Analyzer](#)..... 386
- [Analyzer CLI commands](#)..... 386
- [Client tools](#).....392

About Unisphere Analyzer

Analyzer works within storage domains that the web-based Unisphere software administers. It gathers system performance statistics and presents them in various types of charts. These charts can help you find and anticipate bottlenecks in the disk storage component of a computer system.

Analyzer has two parts: a provider program that accumulates performance information and a user interface (UI). For information about the Analyzer UI, see the Unisphere online help.

The Analyzer enabler must be installed on any system you want to analyze. It controls, accumulates, and manages access to performance information on storage processors (SPs), logical units (LUNs), and disks. The new information displaces the old. This manual uses the term *LUN* to refer to LUNs, metaLUNs, fully provisioned LUNs, and thin LUNs, except where explicitly noted.

Analyzer can display the performance data directly (in real time) or as a file containing past performance data from an archive. It can display real time and archive file data simultaneously. You can compare charts from different systems (using real time data) or you can compare different times on one system (using real time and archive data).

Analyzer chart types

Analyzer has seven types of charts:

- Performance Overview
- Performance Survey
- Performance Summary
- Performance Detail
- IO Size Distribution Summary (for LUNs and metaLUNs only)
- IO Size Distribution Detail (for LUNs and metaLUNs only)
- LUN IO Disk Detail (for LUNs only)

Prerequisites

Analyzer requires the following:

- Unisphere software must be installed and enabled on at least one of the systems in the system domain that you want to analyze.
- Analyzer must be enabled on the system you want to analyze.
- The correct version of the Java Runtime Environment (JRE) must be on the client from which you run the browser.

For the latest requirements, see the release notes supplied with the software.

Analyzer CLI commands

This section explains the `naviseccli` command and the commands for monitoring and measuring the performance of VNX series systems with the Analyzer option.

Note

The commands in this section function only with a system that has the optional Analyzer software installed.

analyzer -set

Modifies the performance logging properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. If one of the switches is specified, the user account needs the administrator or manager role.

DESCRIPTION

The `naviseccli analyzer` command with the `-set` switch lets you modify the performance logging properties.

`analyzer -set` is used with `naviseccli` (see [naviseccli on page 20](#)).

SYNTAX

```
analyzer -set [-narinterval seconds] [-rtinterval seconds]
[-nonstop|-logperiod days]
[-periodicarchiving 0|1] [-default]
```

OPTIONS

`-narinterval seconds`

Changes the polling interval for performance logging. The allowable range is 60 to 3600, and the default is 120. Modifying this value affects the duration of the data in the file. For example, at 600 seconds an archive should have approximately 26 hours of data. If you lower the rate to 1200 seconds, the archive file would have approximately 52 hours of data.

`-rtinterval seconds`

Changes the polling interval for real-time chart windows. The valid range is 60 to 3600; the default is 60.

`-nonstop`

Sets performance logging to run non-stop.

`-logperiod days`

Sets performance logging to run for a period of days. The valid range is 1 to 30 days; the default is 7 days.

`-periodicarchiving 0|1`

When you set `-periodicarchiving` to 1, performance logging automatically creates archive files at periods of 156 samples. The default is 0, no periodic archiving.

`-default`

Resets the values of all performance properties back to their default values. It does not change the `-nonstop` or `-logperiod` switch when performance logging is running.

You cannot specify `-nonstop` and `-logperiod` in the same command line. If you do not specify any switches, you will get back the current values.

EXAMPLE #1

This command sets the archive interval to 20 minutes.

```
naviseccli -h ssl_spa analyzer -set -narinterval 1200
```

EXAMPLE #2

This command sets the archive interval back to 2 minutes.

```
naviseccli -h ss1_spa analyzer -set -default
```

OUTPUT

None if the command succeeds; status or error information if it fails.

analyzer -get

Returns the current values of the performance logging properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli analyzer` command with the `-get` switch lets you view the current values of the performance logging properties for the switch you specify.

`analyzer -get` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -get [-narinterval ] [-rtinterval ] [-logperiod ]
[-periodicarchiving ]
```

OPTIONS

`-narinterval`

Returns the current archive polling interval in seconds.

`-rtinterval`

Returns the current real-time polling interval in seconds.

`-logperiod`

Returns the current logging period (number of days or nonstop).

`-periodicarchiving`

Returns the `-periodicarchiving` setting, No or Yes.

Note

If you do not specify any switches, the `analyzer -get` command displays all the current values.

EXAMPLE

This command returns the current archive interval.

```
naviseccli -h ss1_spa analyzer -get -narinterval
```

OUTPUT

For this example, the output is:

```
Archive Poll Interval (sec): 600
```

analyzer -logging

Resets performance logging.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. You need administrator or manager privileges to use this command.

DESCRIPTION

The `naviseccli analyzer` command with the `-logging` switch lets you reset performance logging.

`analyzer -logging` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -logging [-reset]
```

OPTIONS

`-reset`

Stops performance logging. It deletes all the data that the Analyzer provider has collected. Then it starts performance logging.

EXAMPLE

This command resets performance logging.

```
naviseccli -h ss1_spa analyzer -logging -reset
```

OUTPUT

None if the command succeeds; status or error information if it fails.

analyzer -start

Starts performance logging.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. You need administrator or manager privileges to use this command.

DESCRIPTION

The `naviseccli analyzer` command with the `-start` switch starts performance logging. It turns on statistics logging if it is not on already.

`analyzer -start` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -start
```

EXAMPLE

This command starts performance logging.

```
naviseccli -h ss1_spa analyzer -start
```

OUTPUT

None if the command succeeds; status or error information if it fails.

analyzer -stop

Stops performance logging.

PREREQUISITES

You must have a user account on the system on which you want to execute the command. You need administrator or manager privileges to use this command.

DESCRIPTION

The `naviseccli analyzer` command with the `-stop` switch stops performance logging. It restores the statistics logging state to what it was before logging was started.

`analyzer -stop` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -stop
```

EXAMPLE

This command stops performance logging.

```
naviseccli -h ssl_spa analyzer -stop
```

OUTPUT

None if the command succeeds; status or error information if it fails.

analyzer -status

Displays performance logging status.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli analyzer` command with the `-status` switch lets you see the current status of performance logging. Possible status includes the following:

- Running. Started on date time
- Stopped
- Stopped on date time

`analyzer -status` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -status
```

EXAMPLE

This command provides the current status of performance logging.

```
naviseccli -h ssl_spa analyzer -status
```

OUTPUT

For this example, the output is:

```
Running. Started on 10/17/2006 13:27:32
```

analyzer -archive

Manages archive files.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli analyzer` command with the `-archive` switch lets you list, retrieve, or delete one or more archive files from an SP. It retrieves both the encrypted and decrypted archive files.

SYNTAX

```
analyzer -archive
[[-list][-path pathname][-o][-all|-file filenames][-delete]]
or
analyzer -archive
[-new|-statusnew filename]
```

OPTIONS

- list
Lists all the archive files, ignoring all other switches.
- path *pathname*
Retrieves files to the path you specify. If you do not specify a pathname, files are retrieved to the current directory.
- o
Does not display warning messages.
- all
Retrieves all archive files.
- file *filenames*
Retrieves archive files you specify.
- delete
Deletes files you select or specify with the `-all` or `-file` switches. If you do not specify the `-o` switch, a warning message lists files to delete and asks for confirmation.
- new
Immediately creates a new archive file and output name for the new archive file. If you cannot create a new archive file because less than ten samples were collected, this switch returns the name of the newest archive file.
- statusnew *filename*
Returns the status of the new archive file, *filename*, which you created with the `-new` switch. The possible status is done or not done.

EXAMPLE

This command lists all the archive files.

```
naviseccli -h ssl_spa_analyzer -archive -list
```

OUTPUT

```

Index Size in KB Last Modified Filename
0      39      10/04/2006 16:07:24
APM00035101236_SPA_2006-10-04_20-07-21-GMT_M05-00.nar
1      60      10/06/2006 10:39:56
APM00035101236_SPA_2006-10-06_14-39-56-GMT_M05-00.nar
2     114      10/06/2006 13:49:51
APM00035101236_SPA_2006-10-06_17-49-50-GMT_M05-00.nar
3      68      10/08/2006 12:12:42
APM00035101236_SPA_2006-10-08_16-12-42-GMT_M05-00.nar
4      68      10/08/2006 14:48:42
APM00035101236_SPA_2006-10-08_18-48-42-GMT_M05-00.nar
5      70      10/08/2006 17:24:42
APM00035101236_SPA_2006-10-08_21-24-42-GMT_M05-00.nar
6      70      10/08/2006 20:00:42
APM00035101236_SPA_2006-10-09_00-00-42-GMT_M05-00.nar
7      71      10/08/2006 22:36:42
APM00035101236_SPA_2006-10-09_02-36-42-GMT_M05-00.nar
8      70      10/09/2006 01:12:42
APM00035101236_SPA_2006-10-09_05-12-42-GMT_M05-00.nar
9      70      10/09/2006 03:48:42
APM00035101236_SPA_2006-10-09_07-48-42-GMT_M05-00.nar
10     71      10/09/2006 06:24:42
APM00035101236_SPA_2006-10-09_10-24-41-GMT_M05-00.nar
11     69      10/09/2006 09:00:42
APM00035101236_SPA_2006-10-09_13-00-42-GMT_M05-00.nar
12     72      10/09/2006 11:36:42
APM00035101236_SPA_2006-10-09_15-36-42-GMT_M05-00.nar

```

Client tools

This section describes the Analyzer client tools that you use with `naviseccli`.

analyzer -archivedump

Dumps Analyzer archive files.

PREREQUISITES

The `host`, `username/password` and `scope` options of `naviseccli` are ignored since this command runs locally.

DESCRIPTION

The `naviseccli analyzer` command with the `-archivedump` command dumps the performance data, relationship, and configuration information in a format that you specify.

Note

The `-archivedump` command recognizes and reports an error message when you try to open an encrypted archive file.

The `-archivedump` command includes the following subcommands:

- `-data`
- `-stats`
- `-rel`
- `-config`

`analyzer -archivedump` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -archivedump -data filename(s) [-out filename(s)]
[-join]
[-overwrite y|n] [-xml|-delim code] [-enquote code]
[-eol code]
[-header y|n] [-object codes] [-format codes]
[-stime starttime] [-ftime endtime] [-timezone code]
[-progress]
```

```
analyzer -archivedump -stats filename(s) [-out filename(s)]
[-join]
[-overwrite y|n] [-timezone code]
```

```
analyzer -archivedump -rel filename(s) [-out filename(s)]
[-join]
[-overwrite y|n] [-xml|-delim code] [-enquote code]
[-eol code] [-root objcode] [-level depth] [-progress]
```

```
analyzer -archivedump -config filename(s) [-out filename(s)]
[-join]
[-overwrite y|n] [-xml|-delim code] [-enquote code]
[-eol code] [-object objcode] [-progress]
```


OPTIONS

`-data filename(s)`

Dumps data from archive files. You must provide the archive filename following the `-data` command. The optional switches are as follows:

`-out filename(s)`

Dumps multiple files. Specify multiple output filenames separated by commas. The output filenames list corresponds to the dumped filename list.

Note

You can dump multiple files into a single file using the `-join` switch.

`-join`

Appends the output from multiple files into a single file. The header will be dumped only once if this switch is used and if the header switch is enabled. This switch is valid only if you specify multiple files. If you use this switch with the single file following `-data`, the switch is ignored.

`overwrite y|n`

Works with the `-out` switch. If you use the `-overwrite` switch without the `-out` switch, it is ignored. The `-overwrite` switch applies to all the output files if you specify multiple files to be dumped.

y

Without a warning, overwrites the existing file specified by the `-out` switch.

n

Displays an error message if the file that the `-out` switch specified already exists.

Note

If you try to overwrite a read-only file, an error message appears.

`-xml`

Specifies the output in XML format. This switch is optional, and you cannot use it with the `-delim` type. If you do not specify the `-xml` switch, delimited output results. If you use it with the `-data` command and the `-format` switch, an error message appears.

`-delim code`

Specifies a delimiter that separates the output data. With this switch, you can specify the output as delimited type. It is optional and cannot be used with the `-xml` switch. Specify the delimiter when you use this switch. If you do not specify either `-xml` or `-delim`, delimited output results.

Table 32 Delimiter codes

Delimiter	Code
Comma (default)	cm
Semicolon	sc
Colon	cl
Tab	tb

Table 32 Delimiter codes (continued)

Delimiter	Code
Slash	sl
Backslash	bs
Period	pd
White space	ws

Note

If you use a slash or a period as a delimiter, you must also use the `-enquote` switch with an option other than None.

`-enquote code`

Encloses each performance characteristic within a set of special characters. By default, the special characters None are used. The possible special characters for the switch are listed below:

Table 33 Valid `-enquote` switch special character codes

Enclosing character	Code
None	no
Quote	qu
Apostrophe	ap
Accent character	ac

Note

Use this switch only when you specify `-delimited` as an output type.

`-eol code`

Specifies a particular end-of-line character. The output reaches an end of line when all performance characteristics for one specific object at a particular time have been output. The underlying client system determines the default end-of-line character. The possible end-of-line characters for the switch are listed below:

Table 34 Valid end-of-line characters

End-of-line character	Code
Carriage return	cr
Line feed	lf

Note

Use this switch only when you specify `-delimited` as an output type.

`-header y | n`

Specifies whether to output a column header. This switch is optional. If you do not specify the switch, by default, it outputs a header.

y

Output a header.

n

Do not output a header.

`-object codes`

Specifies the objects for which to collect performance statistics and the order in which to output them. If you do not specify this switch, all object types are dumped in the following order: SPs, public classic LUNs, metaLUNs, fully provisioned LUNs, thin LUNs, Snapshot mount point, private classic LUNs, disks, MirrorView/A, Snap sessions, RAID groups, pools, Snapshots, Consistency Groups, and ports. If the objects have no performance data, no information is dumped.

Table 35 Valid object type codes

Object type	Code
SP	s
LUN ^a	l
MetaLUN	ml
Disk	d
Asynchronous mirror	am
Snap session	ss
RAID group	rg
Port	p
Private classic LUNs ^b	pl
Host LUN ^c	hl
Thin LUN ^d	tl
All LUNs	al
Public/private pool LUNs	plu
Thin pool	tp
Pool ^e	pool
Snapshot ^f	snap
Consistency Group ^g	cg
Snapshot Set ^h	cgsnap
Snapshot mount point ⁱ	sl

- a. RAID group based public or private LUNs
- b. Private LUNs are metaLUN components and other reserved LUNs.
- c. Host LUNs are all LUNs with host I/O.
- d. Thin LUNs are LUNs whose storage capacity grows by using a shared pool of storage. They dump only the host statistics.

Table 35 Valid object type codes (continued)

- e. Thin pool and pool mean the same object. Thin pool is kept for backward compatibility purposes in Unisphere version 1.0
- f. New snapshot object for VNX snapshots.
- g. New object for VNX snapshots.
- h. Snap of a consistency group. New objects for VNX snapshots.
- i. Snapshot mount point for VNX snapshots.

-format codes

Specifies which performance characteristics to output and the order in which they appear. The `-format` switch requires one or more of the codes listed below. If you do not specify this switch, all performance characteristics are output in the order listed below:

Table 36 Performance characteristics without -format code

Performance characteristic	Object type	Code
Object Name	All objects	on
Poll Time	All objects	pt
Owner Array Name	All objects	oan
Current Owner	LUN, metaLUN, pool-based LUN	co
Utilization (%)	LUN, metaLUN, pool-based LUN, disk, SP	u
Queue Length	LUN, metaLUN, pool-based LUN, disk, SP	ql
Response Time (ms)	LUN, metaLUN, pool-based LUN, disk, SP	rt
Total Bandwidth (MB/s)	LUN, metaLUN, pool-based LUN, disk, SP, asynchronous mirror, port	tb
Total Throughput (IO/s)	LUN, metaLUN, pool-based LUN, disk, SP, asynchronous mirror, port	tt
Read Bandwidth (MB/s)	LUN, metaLUN, pool-based LUN, disk, SP, port	rb
Read Size (KB)	LUN, metaLUN, pool-based LUN, disk, SP, port	rs
Read Throughput (IO/s)	LUN, metaLUN, pool-based LUN, disk, SP, port	rio
Write Bandwidth (MB/s)	LUN, metaLUN, pool-based LUN, disk, SP, port	wb
Write Size (KB)	LUN, metaLUN, pool-based LUN, disk, SP, port	ws
Write Throughput (IO/s)	LUN, metaLUN, pool-based LUN, disk, SP, port	wio
Full Stripe Write/s	LUN, metaLUN	fsw

Table 36 Performance characteristics without -format code (continued)

Performance characteristic	Object type	Code
SP Cache Prefetch Bandwidth (MB/s)	LUN, metaLUN	pb
SP Cache Used Prefetches (%)	LUN, metaLUN	up
SP Cache Read Hits/s	LUN, metaLUN	rch
SP Cache Read Misses/s	LUN, metaLUN	rcm
SP Cache Read Hit Ratio	LUN, metaLUN	rchr
SP Cache Write Hits/s	LUN, metaLUN	wch
SP Cache Write Misses	LUN, metaLUN	wcm
SP Cache Write Hit Ratio	LUN, metaLUN	wchr
SP Cache Flush Requests/s	SP	spcfr
Complete histogram (r512b through t512kb)	LUN, metaLUN	histo
Read histogram (r512b through r512kb)	LUN, metaLUN	rhisto
Write histogram (w512b through w512kb)	LUN, metaLUN	whisto
Total histogram (t512b through t512kb)	LUN, metaLUN	thisto
SP Cache Forced Flushes/s	LUN, metaLUN	ff
Disk Crossings (%)	LUN, metaLUN	dcp
Average Busy Queue Length	LUN, metaLUN, pool-based LUN, disk, SP	abql
Disk Crossing/s	LUN, metaLUN	dc
Service Time (ms)	LUN, metaLUN, pool-based LUN, disk, SP	st
LUN Read Crossings/s	metaLUN	krc
LUN Write Crossings/s	metaLUN	lwc
Average Seek Distance (GB)	disk	asd
SP Cache Dirty Pages (%)	SP	dp
SP Cache MBs Flushed (MB/s)	SP	mbf
High Water Flush On	SP	hwfo
Idle Flush On	SP	ifo
Low Water Flush Off	SP	lwfo
Write Cach Flushes/s	SP	wcf
Reads from Snapshot Cache	Snap Session	rfsc
Reads from Snapshot Copy LUN	Snap Session	rfsnl
Reads from Snapshot Source LUN	Snap Session	rfssl
Writes To SnapShot Source LUN	Snap Session	wssl

Table 36 Performance characteristics without -format code (continued)

Performance characteristic	Object type	Code
Writes To Snapshot Cache	Snap Session	wsc
Writes Larger Than Snapshot Cache Chunk Size	Snap Session	wltsccs
Chunks Used in Snapshot Copy Session	Snap Session	cuscs
Average Transfer Size (KB)	Asynchronous Mirror	amats
Time Lag (min)	Asynchronous Mirror	amtl
Data Lag (MB)	Asynchronous Mirror	amml
Cycle Count	Asynchronous Mirror	amcc
Average Cycle Time (min)	Asynchronous Mirror	amact
Read 512B	LUN, metaLUN	r512b
Read 1KB+	LUN, metaLUN	r1kb
Read 2KB+	LUN, metaLUN	r2kb
Read 4KB+	LUN, metaLUN	r4kb
Read 8KB+	LUN, metaLUN	r8kb
Read 16KB+	LUN, metaLUN	r16kb
Read 32KB+	LUN, metaLUN	r32kb
Read 64KB+	LUN, metaLUN	r64kb
Read 128KB+	LUN, metaLUN	r128kb
Read 256KB+	LUN, metaLUN	r256kb
Read 512KB+	LUN, metaLUN	r512kb
Write 512B	LUN, metaLUN	w512b
Write 1KB+	LUN, metaLUN	w1kb
Write 2KB+	LUN, metaLUN	w2kb
Write 4KB+	LUN, metaLUN	w4kb
Write 8KB+	LUN, metaLUN	w8kb
Write 16KB+	LUN, metaLUN	w16kb
Write 32KB+	LUN, metaLUN	w32kb
Write 64KB+	LUN, metaLUN	w64kb
Write 128KB+	LUN, metaLUN	w128kb
Write 256KB+	LUN, metaLUN	w256kb
Write 512KB+	LUN, metaLUN	w512kb
Total 512B	LUN, metaLUN	t512b
Total 1KB+	LUN, metaLUN	t1kb

Table 36 Performance characteristics without -format code (continued)

Performance characteristic	Object type	Code
Total 2KB+	LUN, metaLUN	t2kb
Total 4KB+	LUN, metaLUN	t4kb
Total 8KB+	LUN, metaLUN	t8kb
Total 16KB+	LUN, metaLUN	t16kb
Total 32KB+	LUN, metaLUN	t32kb
Total 64KB+	LUN, metaLUN	t64kb
Total 128KB+	LUN, metaLUN	t128kb
Total 256KB+	LUN, metaLUN	t256kb
Total 512KB+	LUN, metaLUN	t512kb
Queue Full Count	port	qfc
FAST Cache Read Hit Ratio	LUN, Pool	fcrhr
FAST Cache Write Hit Ratio	LUN, Pool	fcwhr
FAST Cache Read Hits/s	LUN, Pool	fcrh
FAST Cache Write Hits/s	LUN, Pool	fcwh
FAST Cache Read Misses/s	LUN, Pool	fcrm
FAST Cache Write Misses/s	LUN, Pool	fcwm
FAST Cache Dirty Pages [%]	SP	fcdp
FAST Cache Flushed MB/s	SP	fcmbf

The options of new performance characteristics are as shown below:

Table 37 Performance characteristics

Optimal performance characteristics	Nonoptimal performance characteristics	Optimal/nonoptimal code	Object type
Utilization-Optimal[%]	Utilization-Nonoptimal[%]	uo/un	LUN, metaLUN, pool-based LUN
Queue Length-Optimal	Queue Length-Nonoptimal	qlo/qln	LUN, metaLUN, pool-based LUN
Response time-Optimal [ms]	Response time-Nonoptimal [ms]	rto/rtn	LUN, metaLUN, pool-based LUN
Total Bandwidth-Optimal [MB/s]	Total Bandwidth-Nonoptimal [MB/s]	tbo/tbn	LUN, metaLUN, pool-based LUN
Total Throughput-Optimal [IO/s]	Total Throughput-Nonoptimal [IO/s]	tto/ttn	LUN, metaLUN, pool-based LUN
Read Bandwidth-Optimal[MB/s]	Read Bandwidth-Nonoptimal[MB/s]	rbo/rbn	LUN, metaLUN, pool-based LUN

Table 37 Performance characteristics (continued)

Optimal performance characteristics	Nonoptimal performance characteristics	Optimal/nonoptimal code	Object type
Read Size-Optimal[KB]	Read Size-Nonoptimal[KB]	rso/rsn	LUN, metaLUN, pool-based LUN
Read Throughput-Optimal[IO/s]	Read Throughput-Nonoptimal[IO/s]	rioo/rion	LUN, metaLUN, pool-based LUN
Write Bandwidth-Optimal[MB/s]	Write Bandwidth-Nonoptimal[MB/s]	wbo/wbn	LUN, metaLUN, pool-based LUN
Write Size-Optimal [KB]	Write Size-Nonoptimal [KB]	wso/wsn	LUN, metaLUN, pool-based LUN
Write Throughput-Optimal[IO/s]	Write Throughput-Nonoptimal[IO/s]	wioo/wion	LUN, metaLUN, pool-based LUN
Average Busy Queue Length-optimal	Average Busy Queue Length-Nonoptimal	abqlo/abqln	LUN, metaLUN, pool-based LUN
Service Time - Optimal [ms]	Service Time - Nonoptimal [ms]	sto/stn	LUN, metaLUN, pool-based LUN
Explicit Trespass Count		etc	LUN, metaLUN, pool-based LUN
Implicit Trespass Count		itc	LUN, metaLUN, pool-based LUN
Implicit Trespass Count-Optimal	Implicit Trespass Count-Nonoptimal	itco/itcn	LUN, metaLUN, pool-based LUN
Explicit Trespass Count-Optimal	Explicit Trespass Count-Nonoptimal	etco/etcn	LUN, metaLUN, pool-based LUN

The following table is an example of the `-archivedump` command, which displays optimal and nonoptimal statistics and also displays the current owner:

Table 38 -archivedump command statistics

Object name	Owner array name	Current owner	Utilization(%)	Utilization-optimal [%]	Response time-optimal [ms]
LUN 2	HK192200361	A	0	0	0
LUN 1	HK192200361	B	8.153079	8.153079	34.986851

`-stime starttime`

Specifies an archive poll time from which to start iterating through the data. The `-stime` switch requires start time to be in the following format:

mm/dd/yyyy hh:mm:ss

where:

mm
is the month of a year (1-12)

dd
is the day of a month (1-31)

yyyy
is the four-digit year

hh
is the hour of the day (0-23)

mm
is the minute in an hour (0-59)

ss
is the second in a minute (0-59)

You must enclose the time in double quotation marks.

The `-stime` switch works in combination with the `-ftime` switch, described below. If the start time is greater than the end time, an error message appears. If you do not specify an end time, this switch defaults to the latest poll time in the archive file.

Note

If you use the `-stime` switch with the `-timezone` switch, you must apply the time zone you specify to the start time.

If the start time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the start time format is not valid, an error message appears.

`-ftime` *endtime*

Specifies an archive poll time at which to stop iterating through the data. The `-ftime` switch requires end time to be in the following format:

mm/dd/yyyy hh:mm:ss

where:

mm
is the month of a year (1-12)

dd
is the day of a month (1-31)

yyyy
is the four-digit year

hh
is the hour of the day (0-23)

mm
is the minute in an hour (0-59)

ss
is the second in a minute (0-59)

You must enclose the time in double quotation marks.

The `-ftime` switch works in combination with the `-stime` switch, described above. If the start time is greater than the end time, an error message appears. If you do not specify an end time, this switch defaults to the latest poll time in the archive file.

Note

If you use the `-ftime` switch with the `-timezone` switch, you must apply the time zone you specify to the end time.

If the end time you specify does not exist in the time interval of the archive file, an error message containing the valid time range of the archive file appears. Also, if the end time format is not valid, an error message appears.

`-timezone code`

Applies a particular time zone to the performance characteristic, Poll Time, included in the output data. By default, Analyzer displays this characteristic in the local time of the client system. The possible time zones for the switch are listed below:

Table 39 Time zone values

Time zone	Code
Local time (default)	local
Universal time	utc

`-progress`

Displays information about the current dump status, including the name of the file being dumped and the percentage dumped. An approximate value is calculated based on the number of objects for the percentage dumped information. This switch works with the `-out` switch. If you use the `-progress` switch without the `-out` switch, it is ignored. Do not use this switch in scripts.

`-stats`

Dumps the statistical information for multiple archive files. Dumped information of different files is separated by a blank line. The output information includes the start time, end time, total number of objects, LUNs, SPs, disks, metaLUNs, pools, thin LUNs, snap sessions, and the number of other objects such as host and system. You can use the following switches with the `-stats` command: `-out`, `-join`, `-overwrite`, and `-timezone codes` (described above). The `-stats` command does not support `-xml` format.

`-rel`

Dumps the relationship information among the various objects in the archive file. You must provide the archive filename following the `-rel` command. You can specify multiple archive files separated by commas to dump the relationship information of all of them. In addition to the switches `-join`, `-out`, `-overwrite`, `-progress`, `-enquote`, `-eol`, `-xml`, `-delim` and `-progress` (described above), you can also use the following optional switches with the `-rel` command:

`-root objcode`

Dumps the relationship information about the specified root and its subnodes. The following object codes specify the object types of a root node:

Table 40 Object types of a root node

Object type	Code
SP	s
Disk	d
Asynchronous mirror	am
Snap session	ss
RAID group	rg
Host	h
Storage System	stor
MetaLUN	ml
LUN ^a	l
port	p
Private classic LUNs (CX4 series only) ^b	pl
Host LUN (CX4 series only) ^c	hl
Thin LUN ^d	tl
Thin pool ^e	tp
All LUNs	al
Public/private pool LUNs	plu
Pool ^f	pool
Snapshot mount point ^g	sl
Snapshot ^h	snap
Consistency Group ⁱ	cg
Snapshot Set ^j	cgsnap

- a. RAID group based public or private LUNs
- b. Private LUNs are metaLUN components and other reserved LUNs.
- c. Host LUNs are all LUNs with host I/O.
- d. Thin LUNs are LUNs whose storage capacity grows by using a shared pool of storage.
- e. Thin pools are sets of disks, all with the same redundancy, that share their user capacity with one or more thin LUNs.
- f. Thin pool and pool mean the same object. Thin pool is kept for backward compatibility purposes in Unisphere version 1.0
- g. New object for MR.
- h. Advanced Snap. New object for MR.
- i. New object for MR.
- j. Snap of consistency group. New object for MR.

Note

Not specifying the switch dumps the relationship information for the system. If you specify multiple switches for this option, an error message appears.

-level *depth*

Specifies the depth of the relationship information to be dumped. Specify the level number as an integer. For example, specifying the level as 2 dumps the

immediate subnodes of all the objects of the type you specified with the `-root` switch. Specifying the level as 1 dumps the name of all object types you specified with the `-root` switch. If you specify a level as 0, an error message appears. Not specifying this switch dumps the complete hierarchy of all the objects of the type specified with the `-root` switch.

`-config`

Dumps the configuration information of objects in the archive file. You must provide the archive filename following the `-config` command. You can specify multiple archive files separated by commas to dump the configuration information for all of them.

You can use the following switches with the `-config` command: `-out`, `-join`, `-overwrite`, `-xml`, `-delim`, `-enquote`, `-eol`, and `-progress` (described above).

`-object objcode`

Specifies for which objects and in what order you want to dump the configuration information. If you do not specify this switch, all object types are dumped in the following order: SPs, public classic LUNs, metaLUNs, fully provisioned LUNs, thin LUNs, Snapshot mount point, private classic LUNs, disks, MirrorView/A, Snap sessions, RAID groups, pools, hosts, systems, Snapshots, Consistency Groups, and ports.

Table 41 Order of dump configuration information

Object type	Code
SP	s
Disk	d
Snap session	ss
Asynchronous mirror	am
RAID group	rg
Host	h
Storage System	stor
metaLUN	ml
LUN	l
Port	p
Private classic LUNs (CX4 series only)	pl
Host LUN (CX4 series only)	hl
Thin LUN	tl
Thin pool	tp
All LUNs	al
Public/private pool LUNs	plu
Pool	pool

EXAMPLE #1

This command lets you dump data from the archive file `archive.nar` without outputting a header.

```
naviseccli analyzer -archivedump -data archive.nar -header n
```

EXAMPLE #2

This command lets you dump statistical information for multiple archive files, `archive1.nar` and `archive2.nar`.

```
naviseccli analyzer -archivedump -stats archive1.nar,archive2.nar
```

EXAMPLE #3

This command lets you dump relationship information of various objects in the archive file `archive.nar` and outputs it in XML format.

```
naviseccli analyzer -archivedump -rel archive.nar -xml
```

EXAMPLE #4

This command lets you dump configuration information of objects in the archive file `archive.nar`; if the output file `archive.csv` exists, this command overwrites it without a warning.

```
naviseccli analyzer -archivedump -config archive.nar -out
archive.csv -overwrite y
```

EXAMPLE 5

```
naviseccli analyzer -archivedump -data archive.nar -out output.csv
```

OUTPUT

Varies depending upon which switch you use.

analyzer -archiveretrieve

Retrieves Analyzer archive files.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli analyzer` command with the `-archiveretrieve` command lets you create the archive files on the SP, retrieve the archive files from the SP, and delete them from the SP.

`analyzer -archiveretrieve` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -archiveretrieve [-file filename] [-location pathname]
[-overwrite y|n] [-retry times] [-v]
```

OPTIONS

`-file filename`

Specifies the name of the archive file once it is retrieved to the client.

`-location pathname`

Specifies the path location of the archive file once it is retrieved to the client system. If you do not specify a location, Analyzer uses the current working directory.

`-overwrite y|n`

Specifies whether to overwrite an existing archive file on the client system. If the specified archive file already exists and this switch is not used, the command fails.

y = Overwrite existing file.

n = Do not overwrite existing file.

`-retry times`

Specifies the number of times to retry a retrieve before giving up completely. A retry occurs only when the actual move of the archive file from the system to the client system fails.

`-v`

As the retrieve process progresses, the status of each stage appears.

EXAMPLE

This command retrieves an archive to the file `archive.nar` in `c:\Temp`, overwrites the existing file and retries three times.

```
naviseccli -h ssl_spa -user username -password password
-scope scope -code analyzer -archiveretrieve -file archive.nar
-location c:\Temp\ -overwrite y -retry 3
```

analyzer -archivemerge

Merges Analyzer archive files.

PREREQUISITES

The host, username/password and scope options of `naviseccli` are ignored since this command runs locally.

DESCRIPTION

The `naviseccli analyzer` command with the `-archivemerge` command lets you merge two archives from the same SP into a single archive.

Note

The `-archivemerge` command recognizes and reports an error message when you try to open an encrypted archive file.

`analyzer -archivemerge` is used with `naviseccli` (refer to [naviseccli on page 20](#)).

SYNTAX

```
analyzer -archivemerge -data archive1 archive2 [-out
outputarchive] [-overwrite y|n]
```

OPTIONS

`-data archive1 archive2`

Merges the data of two archive files that you specify.

`-out outputarchive`

Specifies an output file to which to write the merged archive. Use quotation marks around the filename if it has spaces. If you do not specify this switch, the output merged file will be in the current directory with the default name (`DD_MM_YY_HH_MM_merged.nar`).

`-overwrite y|n`

Specifies whether to overwrite an existing archive file on the client system.

y = Overwrite existing file

n = Do not overwrite existing file

Note

If you try to overwrite a read-only file, an error message appears.

EXAMPLE

This command merges the files C:\Temp\archive1.nar and C:\Temp\archive2.nar and writes the merged archive in C:\Temp\mergedArchive.nar.

```
naviseccli analyzer -archivemerge -data C:\Temp\archive1.nar  
C:\Temp\archive2.nar -out C:\Temp\mergedArchive.nar
```


CHAPTER 9

EMC VNX MirrorView/A CLI

This chapter introduces the EMC® VNX MirrorView™/Asynchronous software and explains the commands for creating and managing the MirrorView/ Asynchronous remote mirrors and consistency groups. This manual refers to the EMC VNX MirrorView/Asynchronous product as MirrorView /A.

Major topics include:

- [MirrorView/A remote mirroring commands](#)..... 410
- [MirrorView/A consistency group commands](#).....431

MirrorView/A remote mirroring commands

This section explains the `naviseccli` commands and the commands for creating and managing remote mirrors on VNX series systems with the MirrorView/A option. These commands let you use MirrorView/A software to create a byte-for-byte copy of one or more local LUNs connected to a distant system server.

Note

The commands in this section function only with a system that has the optional MirrorView/A software installed.

MirrorView/A operations overview

This section provides an overview of MirrorView/A operations and how to manage the systems using the management software.

Procedure

1. Connect the same management software to both systems and configure the software, so that you can manage both systems.

Note

Optionally, you can perform this step at the secondary site as well.

You must manage both systems, which can be in different domains (see the Unisphere online help).

2. Estimate the size of the reserved LUN pool and the reserved LUNs.
3. Establish a usable, two-way connection between the MirrorView/A systems using the `mirror -enablepath` command (see [mirror -enablepath on page 421](#)).
4. If the primary LUN does not exist, bind it on its server's system. Wait for the LUN to finish binding and add it to the storage group.
5. If the secondary LUN does not exist, create a secondary image LUN.

Note

The secondary LUN can be a different RAID type from the primary, as long as the block size matches.

6. Wait for the secondary LUN to finish binding.
7. From the management workstation, create the remote mirror (see [mirror async - create on page 418](#)).

Next add a secondary image. To add a secondary image, see [mirror async - addimage on page 413](#).

Normally, when you add a secondary image to a mirror, the software synchronizes the secondary image with the primary image.

The software will initially copy all the data from the primary LUN to the secondary LUN, in order to ensure that the two LUNs are identical.

At any time in the previous sequence of steps, you can get remote mirror status with the `mirror -async -list` command.

8. Create a consistency group and add the primary image to it (optional). See [mirror - async -creategroup on page 431](#).

9. If a primary failure occurs, Unisphere reports the failure.

If the primary failure is minor, have the primary fixed and resume mirroring.

If the primary failure is catastrophic, the original client that is managing the system may be unusable and thus unable to report the failure. For such a failure, the administrator at the secondary site must set up a client to manage the system (if not already done), then promote the secondary to primary and take other recovery action as needed. This includes assigning the newly promoted LUN to a storage group, if it is not already the appropriate one.

When you use MirrorView/A on a VMware ESX Server, after you promote the secondary image to a primary image, perform the following steps:

- a. If not already assigned, assign the newly promoted primary image to a storage group on the same or standby ESX Server.
- b. Rescan the bus at the ESX Server level.
- c. If not already created, create a virtual machine (VM) on the same or standby ESX Server. The VM is not powered up when you create it.
- d. Assign the newly promoted primary to the VM.
- e. Power up the VM.

If the VM is created and running and you have not already assigned the newly promoted primary to the VM, perform these steps:

- a. Perform steps a and b above.
- b. If you are running ESX Server 2.5.x, power it down.
- c. To assign the newly promoted primary to the VM, use the virtual center interface for ESX Server 3.x and 2.5.x or the Management User interface for ESX Server 2.5.x.

Note

If you are running ESX Server 3.x, you do not need to power down the VM.

- d. If you are running ESX Server 2.5.x, power up the VM.

Note

If you are running ESX Server 3.x, you do not need to power up the VM.

The primary image (which is now the secondary image) will not be accessible to the primary ESX Server.

Note

For configuration restrictions when using VMFS volumes, go to the E-Lab™ Navigator on EMC Online Support, and under the **PDFs and Guides** tab, open the `VMware ESX server.pdf` file.

10. If access to the secondary storage image fails, the primary system will fracture the remote mirror. If the problem with the secondary is minor (for example, replacing a cable), then the administrator can fix it. Mirroring will recover and resynchronize the image, if the image recovery policy is Automatic.

Whenever you want to stop mirroring, you can first fracture and remove the secondary images, and then destroy the mirror. This does not affect any data on either image LUN, and access to the primary LUN is also unaffected. The LUN that held the secondary image now becomes accessible as a regular LUN.

MirrorView/A CLI functions

The VNX CLI for Block provides one of the two interfaces to the MirrorView/A mirroring software. The Unisphere UI provides the other.

You can use the `naviseccli mirror` CLI command to set up and manage remote mirror sessions.

The `naviseccli mirror` command functions are summarized in [naviseccli MirrorView/A command functions on page 412](#).

Table 42 `naviseccli` MirrorView/A command functions

Essential functions (in order performed)	
<code>mirror -enablepath</code>	Opens a path between the local and remote system. Required regardless of the type of connection between systems.
<code>mirror -async -create</code>	Creates a new remote mirror using an existing LUN. The LUN on which the <code>-create</code> command is executed becomes the primary image.
<code>mirror -async -addimage</code>	Adds a secondary image to a mirror using a LUN on a remote system. The default is to synchronize unless you specify otherwise.
<code>mirror -async -changeimage</code>	Changes secondary image properties such as auto recovery setting and recovery policy.
<code>mirror -async -fractureimage</code>	Fractures a secondary image, suspending mirroring of data to that image. A fracture log is maintained, so that the image can be quickly resynchronized later.
<code>mirror -async -removeimage</code>	Removes a secondary image from a mirror.
<code>mirror -async -promoteimage</code>	Promotes a secondary image to primary, while simultaneously demoting the primary image (if still accessible) to secondary. Required in order to access the data in the secondary image in the event of a failure of the primary.
<code>mirror -async -syncimage</code>	Starts synchronization of the secondary image with the primary. Needed if automatic recovery is not on. I/O is allowed with the primary while synchronization occurs. You can specify a synchronization rate to avoid serious performance impact.
<code>mirror -async -info</code>	Displays MirrorView/A information about a system.
<code>mirror -async -list</code>	Displays information on existing mirrors.

Table 43 navisecli MirrorView/A optional functions

Optional reconfiguration functions (alphabetically)	
<code>mirror -async -change</code>	Changes remote mirror properties.
<code>mirror -async -destroy</code>	Terminates remote mirroring of a LUN. This does not affect the data in the LUN, and the server continues to access it.
<code>mirror -disablepath</code>	Severs all paths between the local system SPs and the remote system SPs.

mirror -async -addimage

Adds an image to a previously created mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `navisecli` (see [navisecli on page 20](#)) to add a secondary image to a mirror using a LUN on a remote system. The `-addimage` function is required if you want mirroring to occur. You must direct the command to the system that holds the primary image. This command lets you create primary and secondary images on thin LUNs, if all the participating systems support mirroring on thin LUNs.

Note

The Thin Provisioning™ enabler must be installed to create thin LUNs.

You can use the `-info` function to display the UIDs of systems that have images that can be mirrored.

Note

To add a secondary image, you must have a working connection to the remote system, and the remote LUN must be exactly the same size (number of blocks) as the primary image.

Note

In this syntax "or" means that you can specify the switch, `-arrayhost sp-hostname | sp-IP-address -lun lun-number | -lunuid lun-uid` or you can specify the switch, `-arrayuid system-uid -lunuid lun-uid`. The rest of the syntax is common to both.

SYNTAX

```
mirror -async -addimage -name name | -mirroruid mirroruid
-arrayhost sp-hostname | sp-IP-address
-lun lun-number | -lunuid lun-uid
```

or

```
mirror -async -addimage -name name | -mirroruid mirroruid
-arrayuid system-uid -lunuid lun-uid
[-recoverypolicy auto|manual]
```

```
[-syncrate high|medium|low]
[-manualupdate| -delaystart time|-enddelay time]
[-nosyncrequired]
```

OPTIONS

`-name name`

Specifies the name that you gave to the mirror with the `-create` command.

`-mirroruid uid`

Specifies the unique 16-byte primary image ID (World Wide Name) the software assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the [mirror -async -list on page 425](#) function).

`-arrayhost sp-hostname|sp-IP-address`

Specifies the hostname or network address of the remote SP that will hold the secondary image. The secondary image will be added to the remote mirror connected to this host's SP. Use this switch or `-arrayuid`.

`-lun lun-number`

Specifies the LUN number that will serve as the secondary image of the mirror.

`-arrayuid system-uid`

Specifies the unique 8-byte system ID of the remote SP. You can determine this by using the `mirror -info` command or the `getarrayuid` command. For any letters in the ID, use the same case that the software assigned when the mirror was created. Use this switch or `-arrayhost`.

`-lunuid lun-uid`

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to be added as a secondary image. If you specified `-arrayuid`, then `-lunuid` is required. If you specified `-arrayhost`, then you can use either `-lun` or `-lunuid`.

Note

When you specify `-arrayuid` and `-lunuid`, you must add the mirror driver on the remote LUN using the `-setfeature` command before executing the `-addimage` command.

`-recoverypolicy manual|auto`

Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:

`manual`: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

`auto`: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.

`-syncrate high|medium|low`

Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other system operations.

`-manualupdate`

Indicates that you must explicitly update the image.

`-delaystart time`

Specifies the time (in minutes) from the beginning of the previous update to the start of the next update. The current update must complete before the next one can start.

If the update is still in process when the time period expires, the next update will start immediately once the current update completes.

`-enddelay time`

Specifies the time (in minutes) from the end of the previous update to the next update. A value of 0 causes updates to occur as fast as possible.

`-nosyncrequired`

Specifies that the image will be added in a synchronized state; thus, no initial synchronization will occur.

EXAMPLE

Any of the following commands adds an image to a mirror. Note that if you specify the remote system using the `-arrayuid` switch, before issuing the command you must explicitly add drivers to the LUN stack using the `-setfeature` command. For SP `ss1_spa`, any of the above commands adds to the remote mirror `db_mirror` the LUN in system `50:06:06:10:00:FD:A1:69` (or a LUN on the remote system `cpc426`) whose unique LUN ID (world wide name, WWN) is `60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11` or whose LUN number is 9. See also the example in [mirror -async -create on page 418](#).

```
naviseccli -h ss1_spa mirror -async -addimage -name db_mirror
-arrayuid 50:06:06:10:00:FD:A1:69
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

```
naviseccli -h ss1_spa mirror -async -addimage -name db_mirror
-arrayhost cpc426 -lun 9
```

```
naviseccli -h ss1_spa mirror -async -addimage -name db_mirror
-arrayhost cpc426
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -change

Changes mirror properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to change one or more parameters of a previously created mirror.

You must direct the command to the system that holds the primary image. You must specify at least one of the optional switches with this command. If you omit the override option, then the CLI displays a message for each parameter to confirm the change.

SYNTAX

```
mirror -async -change -name name | -mirroruid uid
[-description description]
[-requiredimages num_of_images] [-o]
[-enablereadonlysecondary yes|no]
```

OPTIONS

- name *name*
Specifies the name the `-create` command gave the mirror. You can specify either the mirror name or unique ID (next).
- mirroruid *uid*
Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).
- description *description*
Provides a meaningful title for the mirror (for example, Employee and Benefits Table for Human Resources). The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.
- requiredimages *num_of_images*
Specifies the minimum number of secondary images that should be accessible to the mirror. The default value is 0. Allowed values are 0, 1, and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For the remote mirror `dbmirror` on SP `ss0_spa`, this command reduces the minimum number of accessible secondary images to 0. Use the `-list` command to confirm that the changes were made.

```
naviseccli -h ss0_spa mirror -async -change -name
dbmirror -requiredimages 0
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -changeimage

Changes secondary image properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to change one or more parameters of a mirror image.

You must direct the command to the system that holds the primary image. You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

```
Do you want to change the sync rate from Medium to High (y/n)?
To change a parameter, enter y.
```

SYNTAX

```
mirror -async -changeimage -name name | -mirroruid uid
-imageuid image-uid
```



```
[-recoverypolicy manual|auto]
[[-manualupdate] | [-delaystart time]]
[-enddelay time]
[-syncrate high|medium|low] [-o]
```

OPTIONS

- name *name*
Specifies the name you gave to the remote mirror in the `-create` command. You can specify either the mirror name or uid (next).
- mirroruid *uid*
Specifies the unique 16-byte ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).
- imageuid *image-uid*
Specifies the unique 8-byte ID of the secondary image you want to change. For any letters in the ID, you must specify the same case that the software assigned to the image when the mirror was created (use the `mirror -list` command).
- recoverypolicy manual|auto
Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:
manual: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

auto: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.
- manualupdate
Indicates that you must explicitly update the image.
- delaystart *time*
Specifies the time (in minutes) from the beginning of the previous update to the start of the next update. The current update must complete before the next one can start. If the update is still in process when the time period expires, the next update will start immediately once the current update completes.
- enddelay *time*
Specifies the time (in minutes) from the end of the previous update to the next update. A value of 0 causes updates to occur as fast as possible.
- syncrate high|medium|low
Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other system operations.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For system `server1`, mirror `db_mirror`, image `50:06:06:10:00:FD:A1:6`, this command changes the synchronization rate to high. Use the `-list` command to confirm that the changes were made.

```
naviseccli -h server1 mirror -async -changeimage -name
db_mirror -imageuid 50:06:06:10:00:FD:A1:6 -syncrate high
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -create

Creates a new asynchronous mirror using an available LUN for mirroring.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to create a new mirror using an existing LUN. The LUN on which you run the command becomes the primary LUN. But, for mirroring to occur, you must add an image with the `-addimage` function. Later, you can change some of the parameters you specify here with the `-change` function. This command supports thin LUNs after you install the latest Unisphere version, which contains thin LUN support for MirrorView.

If you specify a minimum number of required images (other than the default value zero) in the `-create` command, the mirror will start in the attention state, since you have not yet added the required secondary images. To avoid this problem, do not specify a minimum number of images with `-create`; instead, create the mirror, add the image, and then if you want a minimum number of required images, change the mirror with the `-change` command.

SYNTAX

```
mirror -async -create -name name
-lun lun_number
[-description description]
[-requiredimages num_of_images]
[-o]
```

OPTIONS

`-name name`

Lets you assign a meaningful name to the mirror; for example "Employee Benefits". To include spaces in the name, enclose the entire string in quotation marks as in the example. You can use this name to manage the mirror. The name can include as many as 35 characters.

`-lun lun_number`

Specifies the LUN number that will serve as the primary image of the mirror you are creating.

`-description description`

Lets you assign a description for the mirror. The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.

`-requiredimages num_of_images`

Specifies the minimum number of secondary images that should be accessible to this mirror. The default value is 0. Allowed values are 0, 1 and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

The `mirror -async -create` command creates the mirror named `db_mirror` on LUN 112. The mirror automatically activates.

```
naviseccli -h ssl_spa mirror -async -create -name db_mirror -lun 112
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -destroy

Destroys a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to destroy a mirror. It is the counterpart of `-create`. It ends mirroring of the LUN that is part of the mirror. This does not affect the data in the LUN, and the server continues to access it.

Note

You cannot destroy a mirror while it is a member of a consistency group.

You must remove the secondary image in the mirror before issuing the command to the system holding the primary image. If you cannot remove the secondary image or access the primary image, you can use the `-force` option. You must direct the command to the system that holds the primary image unless `-force` is specified, in which case, you can direct it to either system.

CAUTION

Use `-force` only if all other options to remove an image and the mirror have failed.

If you omit the override option, a message to confirm the destruction of the mirror appears:

```
Do you really want to destroy the mirror <mirror-name> (y/n)?
To destroy the mirror, answer y; to take no action, enter n.
```

SYNTAX

```
mirror -async -destroy -name name | -mirroruid uid
[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the mirror with the `-create` command. You can also use the remote mirror ID.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

`-force`

Forcibly destroys the mirror. This option affects only the local system.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, this command destroys the remote mirror that was created with the name `db_mirror`.

```
naviseccli -h ssl_spa mirror -async -destroy db_mirror
```

```
Do you really want to destroy the mirror db_mirror?(y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages if the command fails for some reason.

mirror -disablepath

Disables the communication paths created by `-enablepath` to remote mirror SPs.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 20](#)) to sever the path established by `-enablepath` in the remote system.

When issued to any SP, the command severs paths for both SPs in both the primary and secondary systems; that is, it severs these paths if:

- SP A of the local system is connected to SP A of the remote system
- SP A of the remote system is connected to SP A of the local system
- SP B of the local system is connected to SP B of the remote system
- SP B of the remote system is connected to SP B of the local system

This means you can issue the command to either SP of either system to sever the paths between the systems you are targeting.

If you issue the command while a mirror exists, the CLI issues a warning message; if you continue, any mirror connected over the specified path will be fractured.

Note

If both asynchronous and synchronous mirrors are installed, you need to run the command only once. It will take effect for both mirrors.

If you omit the override option, the CLI displays a message to confirm the disabling of the path. To disable the path, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -disablepath connected-SP-hostname [-connectiontype  
fibre|iscsi] [-o]
```

OPTIONS

connected-SP-hostname

Specifies the hostname of SP A or SP B of a connected system to disconnect.

`-connectiontype fibre | iscsi`

Specifies the connection type as fibre or iSCSI for the path you are disabling. If you do not specify the `-connectiontype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

-o

Executes the command without prompting for confirmation. Without this switch, the CLI displays a confirmation message for each SP.

EXAMPLE #1

For SP `ss1_spa`, this command and confirmation remove the paths between the local SPs and the remote mirror SPs.

```
naviseccli -h ss1_spa mirror -disablepath ss27_spa -connectiontype
fibre
```

```
Disabling of the path will stop mirroring of data to all
images if any among these arrays. A system fracture of
all mirrors between two arrays on following storage
processors will occur due to this action.>128.221.39.29>10.14.20.154
Do you want to continue (y/n)? y
```

EXAMPLE #2

The `mirror -info` command displays the updated status.

```
naviseccli -h ss1_spa mirror -async -info
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to stdout.

mirror -enablepath

Establishes a path for remote mirroring between a local and remote system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 20](#)) to create a path between the specified SP and its peer SP to a remote SP and its peer. These paths will serve as paths for remote mirroring. A remote mirror requires a path between the primary and secondary systems.

When issued to any SP, the command creates paths for both SPs in both systems; that is, it creates these paths:

- SP A of the local system to SP A of the other system
- SP A of the other system to SP A of the local system
- SP B of the local system to SP B of the other system
- SP B of the other system to SP B of the local system

The SP As or the SP Bs in both systems must be physically connected.

If you cannot enable paths on one SP due to some error, an informational message appears and the paths are enabled for the other SP. To obtain status information, run the `mirror -info -systems` command.

Note

If both asynchronous and synchronous mirrors are installed, you will need to run the command only once. It will take effect for both mirrors.

SYNTAX

```
mirror -enablepath SP-hostname [-connectiontype fibre|iscsi]
```

OPTIONS

SP-hostname

Specifies the hostname of SP A or SP B of the remote system on which you want to enable remote mirroring.

`-connectiontype fibre | iscsi`

Specifies the connection type as fibre or iSCSI for the path you are enabling. If you do not specify the `-connectiontype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

Note

To change the connection type, you must first disable the current enabled path. To modify, test, or delete connections or change the credentials, see the *EMC Navisphere Command Line Interface (CLI) Reference*.

EXAMPLE

For SP `ss1_spa`, this command enables the paths between `ss1_spa` and its peer SP and `ss27_spa` and its peer SP. The `mirror -info` command that follows displays the updated status.

```
naviseccli -h ss1_spa mirror -enablepath ss27_spa -connectiontype iscsi
```

```
naviseccli -h ss1_spa mirror -async -info
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to stdout.

mirror -async -fractureimage

Fractures an image in a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to start an administrative fracture of a secondary image. It also prevents scheduled updates from starting in the future. You must direct the command to the system that holds the primary image.

(An automatic fracture, called a system fracture, occurs if a secondary image becomes inaccessible to the primary image. After a system fracture, the system periodically tries to communicate with the secondary image and — if it succeeds — starts synchronization based on the fracture log. If you issue the `-fractureimage` command while the mirror is system fractured, the system will not try to communicate or resynchronize until you issue a `-syncimage` command.)

The software displays a message to confirm the fracture:

```
Do you want to fracture the image image UID (y/n)?
```

To fracture the mirror, answer **y**; to take no action, enter **n**.

SYNTAX

```
mirror -async -fractureimage -name name | -mirroruid uid  
-imageuid uid  
[-o]
```

OPTIONS

- name *name*
Specifies the name you gave to the mirror with the `-create` command.
- mirroruid *uid*
Specifies the unique 16-byte mirror ID (World Wide Name) that the software assigned when it created the mirror.
- imageuid *uid*
Specifies the unique 8-byte image ID of the secondary image to be fractured.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa` of the mirror named `db_mirror`, this command fractures the secondary image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -async -fractureimage -name
db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

```
Do you really want to fracture the image (y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -info

Displays general information about asynchronous mirroring in a system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command, when used with `naviseccli` and no other switches, lists all general information for asynchronous mirror support. Issuing this command with one or more options displays the information based on the option issued. See the example for sample output.

SYNTAX

```
mirror -async -info [-cancreate]
[-maxmirrors]
[-maxremoteimages]
[-mirroredluns]
[-mirrorableluns]
[-systems [-connectiontype]]
[-all]
```

OPTIONS

- cancreate
Indicates whether an asynchronous mirror can be created in this system: yes or no.
- maxmirrors
Displays the maximum number of mirrors that can be created in this system.
- maxremoteimages
Specifies the maximum number of secondary images allowed. A maximum of one image is allowed.
- mirroredluns

- Lists the LUN numbers of LUNs that are mirrored in this system. The list includes LUNs used as both secondary and primary images.
- mirrorableluns
 - Lists LUNs that can be mirrored in this system but are not mirrored at present.
- systems
 - Lists remote systems that can be used for a secondary image.
- connectiontype
 - Displays the connection type (fibre or iSCSI) for the enabled mirror. For unexpected error cases, it displays the value, unknown.
- all
 - Lists all properties of a session without having to type in any other switch.

EXAMPLE #1

For SP `ss1_spa`, this command displays potential and current remote mirror information.

```
naviseccli -h ss1_spa mirror -async -info
```

```
Can a mirror be created on this system: YES
Maximum number of remote images: 1 Logical
Units that are mirrored in Primary Images: 128 5
Logical Units that are mirrored in Secondary Images: 0
Logical Units that can be mirrored: 1 2 3 4 10 11 12 13
14 21 22 58 59
Remote systems that can be enabled for mirroring:
Array UID: 50:06:01:60:90:20:27:47
Status: Enabled on both SPs Enabled Connection
Type(s): iSCSI Disabled Connection Type(s): Fibre Maximum
Number of possible Mirrors: 100
```

EXAMPLE #2

The following example shows the output for the `mirror -async -info -all` command.

```
mirror -async -info -all
```

```
Can a mirror be created on this system: YES
Maximum number of Remote Images: 1
Logical Units that are mirrored in Primary Images: 2
Logical Units that are mirrored in Secondary Images:
Logical Units that can be mirrored: 3
Remote systems that can be enabled for mirroring:
Array UID: 50:06:01:60:88:60:01:8C
Remote systems that are enabled for mirroring:
Array UID: 50:06:01:60:88:60:07:4B
Status: Enabled on both SPs
Enabled Connection Type(s): Fibre
Disabled Connection Type(s):
Maximum number of possible Mirrors: 256
Logged-in status of Remote systems:
Array UID: 50:06:01:60:88:60:01:8C
Fibre Logged-in Status: SPA:Y; SPB:Y
Array UID: 50:06:01:60:88:60:07:4B
Fibre Logged-in Status: SPA:Y; SPB:Y
MirrorView port(s):
SP A:
I/O Module Slot: 0
Physical Port ID: 0
Logical Port ID: 0
Port Type: Fibre
SP B:
I/O Module Slot: 0
Physical Port ID: 0
```



```
Logical Port ID: 0
Port Type: Fibre
```

OUTPUT

If the version of the VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`.

mirror -async -list

Displays information about existing asynchronous mirrors.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -async mirror` command with `-list` and no other switches lists all information for all mirrors on the system. Issuing this command with one or more options displays the information based on the option(s) issued. You can direct the command to any system; however, only the system holding the primary image of a mirror will always have the latest information about the mirror. Status obtained from other systems may be stale; do not use it unless the primary system is unavailable. See the example for sample output.

SYNTAX

```
mirror -async -list [-name [[name]] [-mirroruid] [uid]
[-description]
[-faulted]
[-images]
[-imagecount]
[-imagesize]
[-lun]
[-ownergroupuid]
[-ownergroupname]
[-requiredimages]
[-state]
[-timestamp]
[-transition]
```

OPTIONS

- `-name name`
Specifies the name you gave the mirror with the `-create` function. You can use the mirror name or unique ID to obtain information on a mirror.
- `-mirroruid uid`
Specifies the unique 16-byte remote mirror ID (World Wide Name) the software assigned when it created the mirror.
- `-description`
Lists the meaningful title for the mirror, if any. The software does not use the description to identify the mirror.
- `-faulted`
Lists whether the remote mirror is faulted: yes or no. If it is faulted, the mirror is either fractured or in the attention state.
- `-images`
Displays image information, such as the role of the image (primary or secondary), image ID, WWN of the LUN containing the image, recovery policy, sync rate, image state, image condition, and image transitioning. It also displays the update type (Manual, Time elapsed since the start of last update, and Time elapsed since the

- end of the last update), the update period in minutes (displays N/A where not applicable), the time in minutes since the last update, the time in minutes until the next update, the last image error (why the image is faulted; displays the Not Available if no error).
- imagecount
Lists the current number of images that compose the mirror. This count includes the primary image.
- imagesize
Lists the size of the image (capacity of the LUN containing the image) in user blocks.
- lun
Lists the LUN number of the local image.
- ownergroupid
Lists the unique ID value of the consistency group of which the mirror is a member.
- ownergroupname
Lists the unique name of the consistency group of which the mirror is a member.
- requiredimages
Lists the minimum number of secondary images required for the mirror.
- state
Describes the operational state of the mirror. Values are: active and attention.
- timestamp
Returns the clock time when the last update started.
- transition
Specifies whether the mirror is in a changing state.

EXAMPLE

```
naviseccli -h ssl_spa mirror -async -list
```

```
MirrorView Name:  pbc async ml 80
MirrorView Description:
MirrorView UID:
54:33:20:90:60:01:06:50:03:00:00:00:00:00:00:00
Logical Unit Numbers:  80
Remote Mirror Status:  Mirrored
MirrorView State:  Active
MirrorView Faulted:  NO
MirrorView Transitioning:  NO
Minimum number of images required:  0
Image Size:  409600
Image Count:  2
Images:
Image UID:  50:06:01:60:90:20:33:54
Is Image Primary:  YES
Logical Unit UID:  60:06:01:60:2F:E7:0C:00:3C:5B:
04:A2:50:4C:D8:11
Image Condition:  Unknown Image Condition

Image UID:  50:06:01:60:90:20:29:85
Is Image Primary:  NO
Logical Unit UID:  60:06:01:60:3F:E1:0D:00:A4:C9:01:63:B6:4C:D8:11
Image State:  Synchronized
Image Condition:  Normal
Recovery Policy:  Automatic
Synchronization Rate:  Medium
Image Faulted:  NO
Image Transitioning:  NO
Synchronizing Progress(%):  0
Update Type:  Manual
Time in secs since previous update:  12550
Time in secs until next update:  Not Available
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported error` message is printed to `stdout`. You may also receive other error messages.

mirror -async -promoteimage

Promotes a secondary image to primary.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command promotes a secondary image to the role of the primary image. The command is successful if it can promote the secondary to a primary that is server-accessible. You must direct this command to a system that holds the secondary image — never to the system that holds the primary. The command will fail if you issue it to the system that holds the primary image.

You can promote a secondary image if the secondary is in a synchronized or consistent state. There must be no current update, in addition to the state restrictions listed.

When you promote a secondary image to primary, the software assigns a new mirror UID to the promoted image to distinguish it from the old mirror. The new status of the old primary image depends on whether the old primary image is accessible when promotion occurs.

After promoting an image, you must explicitly add the newly promoted image to an appropriate storage group. EMC recommends that you also remove the old primary from its storage group, if possible.

In normal operation when there is no failure, you can (but probably would not want to) promote a secondary image. If you do so, the primary image is demoted to secondary. The new primary image inherits the mirror properties of the demoted primary image.

Before issuing this command consider the following:

- If possible, the secondary image should be synchronized (synchronized state) with the primary image. If the secondary image is in the consistent state (not synchronized state) when you promote it to primary image, data will probably be lost during the promotion process.
The secondary image to be promoted must be in a consistent or synchronized state. If the secondary image is out-of-sync or synchronizing, then you cannot promote it; an error will result.
- You must stop I/O to an accessible primary image before you issue this command. If I/O continues to the original primary, data may be lost as the software demotes the primary to secondary and/or an SP may reboot unexpectedly during the attempted promotion.

SYNTAX

```
mirror -async -promoteimage -name name | -mirroruid uid
-imageuid imageuid
[-type normal|oos|local ] [-o]
[-type type]
[-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-create` command. You can use the mirror name or unique ID to promote a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte mirror ID (World Wide Name) that was assigned when the mirror was created. For any letters in the ID, see comments under `-imageuid` below.

`-imageuid uid`

Specifies the unique 8-byte image ID of the secondary image you want to promote. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the `mirror -list` function).

`-type normal|oos|local`

Checks the status of the mirror images. Normal (the default) directs the driver to perform internal safety checks and return the appropriate error if the promotion would cause an out-of-sync state for the old primary or remote image if the mirror is unreachable. OOS indicates that no safety check is performed, and the promote will occur even if it results in an out-of-sync state for the old primary. Local results in promoting the local image only to the primary. The driver attempts to contact the old primary, and if it reaches it, the old primary removes its secondary image.

`-o`

Runs the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the remote mirror named `db_mirror`, this command promotes the secondary image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -async -promoteimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -removeimage

Removes a secondary image from a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to remove a secondary image from the remote mirror. You must issue this command to a system that holds the primary image.

Note

If an update is in progress, this command will not execute. We recommend that you fracture the secondary image first, and then you can issue the `mirror -async -removeimage` command.

If you omit the `-override` option, the software displays a message to confirm the remove operation:

```
Removing of the image will stop mirroring of data to the
specified image. It may cause the minimum number of images
required parameter to be violated that will cause the mirror to
go into attention state. Do you really want to remove the
image (y/n)?
```

Answering yes will remove the image.

SYNTAX

```
mirror -async -removeimage -name name|-mirroruid mirroruid
-arrayhost sp-hostname|sp-IP-address|-imageuid imageuid
[-o]
```

OPTIONS

-name *name*

Specifies the name you gave the mirror with the `-create` command. You can use the mirror name or unique ID to remove a secondary image.

-mirroruid *uid*

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

-arrayhost *sp-hostname* | *sp-IP-address*

Specifies the hostname or network address of the remote SP that holds the secondary image you want to remove. If you specify this switch, the driver(s) from the secondary LUN stack is removed.

-imageuid *image-uid*

Specifies the unique 8-byte image ID of the secondary image you want to remove. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the `mirror -list` command). If you specify this switch, you must explicitly remove the driver(s) from the secondary LUN stack using the `-setfeature` command.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP address `ss1_spa`, the owner of the mirror named `db_mirror`, this command removes the image with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -removeimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -setfeature

Adds or removes an asynchronous mirror driver to or from the LUN stack.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see `naviseccli`) to add the asynchronous mirror driver to or remove it from the LUN stack depending on whether you specify the `-on` or `-off` option. You must issue this command to the system where the LUN resides.

Note

`-setfeature` is needed only if you want to use the `-arrayuid` switch to specify the remote system for the `-addimage` command.

SYNTAX

```
mirror -async -setfeature -on | -off -lun lun_number | -lunuid uid
```

OPTIONS

-on

Adds the driver to the LUN stack.

-off

Removes the driver from the LUN stack.

-lun *lun_number*

Specifies the LUN number to which you want to add or from which you need to remove the driver.

-lunuid *uid*

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to which you will add the driver or from which you will remove the driver.

EXAMPLE

For SP `ssl_spa`, this command starts adding the driver to LUN 112.

```
naviseccli -h ssl_spa mirror -async -setfeature -on -lun 112
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -syncimage

Starts an update of the secondary image.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 20](#)) to begin the synchronization operation on a secondary image that was previously administratively fractured or on one where the mirror image condition is normal. You must issue this command to the system that holds the primary image. The secondary image should be fractured when performing this operation. If you specify a secondary image that is not fractured, an internal error message will be returned.

The software displays a confirmation message:

```
During synchronization, array performance will be impacted. Do you want to sync the image (y/n)?
```

Answering **y** will start synchronization.

SYNTAX

```
mirror -async -syncimage -name name|-mirroruid mirroruid -imageuid imageuid [-o]
```

OPTIONS

-name *name*

Specifies the name given to the mirror with the `-create` command. You can use the mirror name or unique ID to update a secondary image.

-mirroruid *uid*

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

`-imageuid imageuid`

Specifies the unique 8-byte image ID of the secondary image you will synchronize. For any letters in the ID, you must specify the same case as the software assigned to the image when the software created the mirror (use the `mirror -list` command).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the mirror named `db_mirror`, this command starts synchronizing the image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -async -syncimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

```
During synchronization, array performance will be impacted. Do you
want to sync the image (y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

MirrorView/A consistency group commands

This section explains the commands for creating and managing groups of mirrors on VNX series systems with the EMC MirrorView/A option. This section refers to groups of mirrors as *consistency groups*.

Note

The LUNs in a MirrorView/A consistency group must reside on the same system, unlike the volumes in a Symmetrix consistency group, which can reside on multiple systems.

mirror -async -creategroup

Establishes a consistency group to which you can later add mirrors.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

This command is used with `naviseccli` to create a consistency group to which you can add mirrors with subsequent commands. You must direct the command to the system on which you are creating the MirrorView/A group.

SYNTAX

```
mirror -async -creategroup -name name | -mirroruid uid
[-description description]
[-manualupdate | -delaystart minutes | -enddelay minutes]
[-syncrate rate] [-o]
[-recoverypolicy policy] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group. Use only alphanumeric characters for the consistency group name. The name can be up to 35 characters.

`-description description`

Specifies a textual description of the consistency group.

`-manualupdate`

Specifies that you must explicitly start each update of the group. If you specify neither `-manualupdate`, `-delaystart`, nor `-enddelay`, the default is `-enddelay 60`.

Note

When a consistency group begins an update, the mirrors within the group begin updating and end whenever each individual update finishes. The consistency group update completes when the last mirror completes. If an update is not running, you can manually start an update even if an automatic update interval is set.

`-delaystart minutes`

Specifies the time (in minutes) from the start of the previous update until the start of the next one. If the previous update has not completed, the next update is delayed until the update in progress has completed. If you specify neither `-manualupdate`, `-delaystart`, nor `-enddelay`, the default is `-enddelay 60`.

`-enddelay minutes`

Specifies the time (in minutes) from the end of the previous update until the start of the next one. Specifying a value of 0 results in the group being updated as fast as possible. If you specify neither `-manualupdate`, `-delaystart`, nor `-enddelay`, the default is `-enddelay 60`.

`-syncrate rate`

Specifies a relative value (low, medium, or high) for the priority of completing updates. Values for rate are low, medium, or high. High completes updates faster, but may significantly affect system performance for server I/O requests. Low completes updates slower, but also minimizes the impact on other system operations.

`-recoverypolicy policy`

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

`manual`: Update of the secondary image does not begin until you explicitly issue a `synchronize` command.

`auto`: Update of the secondary image automatically begins as soon as the connection between the two systems is again operational.

`-o`

Executes the command without prompting for confirmation.

Note

All the switches, except `-name` and `-o`, are meaningful only if the system receiving the command hosts the primary images for the mirrors in the consistency group.

EXAMPLE

For SP `ss1_spa`, this command creates a consistency group with the name `dbgroup`.

```
naviseccli -h ss1_spa mirror -async -creategroup -name dbgroup
```


OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -destroygroup

Destroys a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to destroy a consistency group. It is the counterpart of `-creategroup`.

Note

Make sure that the consistency group is empty before destroying it.

If you omit the override option, the system displays a message to confirm the destruction of the consistency group:

```
Do you really want to destroy the group <group-name>? (y/n)?
To destroy the consistency group, answer y; to take no action, enter n.
```

SYNTAX

```
mirror -async -destroygroup -name name | -groupid id
[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-creategroup` command. You can also use the consistency group ID.

`-groupid id`

Specifies the consistency group unique ID assigned when the consistency group was created.

`-force`

Destroys the consistency group even if it is not empty.

Note

Particularly in a failure situation you can use the `-force` switch to destroy a group that still has members.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command destroys the consistency group that was created with the name `dbgroup`.

```
naviseccli -h ss1_spa mirror -async -destroygroup dbgroup
```

```
Do you really want to destroy the group dbgroup? (y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages if the command fails.

mirror -async -addtogroup

Adds a mirror to a previously created consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to add one mirror at a time to a consistency group.

Note

The mirror you specify must have its primary image on the system. No two secondary images of mirrors in the consistency group can be on different systems.

SYNTAX

```
mirror -async -addtogroup -name name | -groupid id
-mirrorname name | -mirroruid id
```

OPTIONS

`-name name`

Specifies the name of the consistency group to which you want to add the mirror.

`-groupid id`

Specifies the unique ID of the consistency group to which you want to add the mirror.

`-mirrorname name`

Specifies the name of the mirror you want to add to the consistency group.

`-mirroruid id`

Specifies the World Wide Name of the mirror you want to add to the consistency group.

EXAMPLE

For SP `ssl_spa`, this command adds the mirror, `Sales Accounts Mirror to Europe`, to the consistency group, `dbgroup`.

```
naviseccli -h ssl_spa mirror -async -addtogroup -name dbgroup
-mirrorname "Sales Accounts Mirror to Europe"
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -removefromgroup

Removes a mirror from a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to remove one mirror at a time from the consistency group.

Note

The mirror you specify must have its primary image on the system. This does not apply if you use the `-force` switch.

If you omit the `-o` option, the software displays a message to confirm the remove operation. Answering yes will remove the mirror.

SYNTAX

```
mirror -async -removefromgroup -name name | -groupid id
-mirrorname name | -mirroruid id[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group from which you want to remove the mirror.

`-groupid id`

Specifies the unique ID of the consistency group to which you want to add the mirror.

`-mirrorname name`

Specifies the name of the mirror you want to remove from the consistency group.

`-mirroruid id`

Specifies the World Wide Name of the mirror you want to remove from the consistency group.

`-force`

Removes the mirror even if the connection with the remote system is not functioning.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command removes the mirror with the unique ID shown.

```
naviseccli -h ssl_spa mirror -async -removefromgroup -name dbgroup
-mirroruid 83:5A:20:90:60:01:06:50:01:00:00:00:00:00:00:00
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -change group

Changes consistency group properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to change one or more parameters of a consistency group.

You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

Do you want to change the group synchronization rate from Medium to High (y/n)?

To change a parameter, enter **y**.

SYNTAX

```
mirror -async -change group -name name | -groupid id
[-description description]
[-manualupdate | -delaystart minutes | -enddelay minutes]
[-syncrate rate]
[-recoverypolicy policy] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to change. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to change.

-description *description*

Specifies a textual description for the consistency group.

-manualupdate

Specifies that you must explicitly start each update of the group.

-delaystart *minutes*

Specifies the time (in minutes) from the start of the previous update until the start of the next one. If the previous update has not completed, the next update is delayed until the update in progress has completed.

-enddelay *minutes*

Specifies the time (in minutes) from the end of the previous update until the start of the next one. Specifying a value of 0 results in the group being updated as fast as possible.

-syncrate *rate*

Specifies a relative value (low, medium, or high) for the priority of completing updates. Values for rate are low, medium, or high. High completes updates faster, but may significantly affect system performance for server I/O requests. Low completes updates slower, but also minimizes the impact on other system operations.

-recoverypolicy *policy*

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

manual: Update of the secondary image does not begin until you explicitly issue a `synchronize` command.

auto: Update of the secondary image automatically begins as soon as the connection between the two systems is again operational.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For system `server1`, consistency group `dbgroup`, this command changes the synchronization rate to `high`.

```
naviseccli -h server1 mirror -async -change group -name dbgroup
-syncrate high
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

Use the `-listgroups` command on the mirror to confirm that the changes were made.

mirror -async -syncgroup

Starts an update of the consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to begin the synchronization operation on a consistency group. All the images in the consistency group must be primary.

SYNTAX

```
mirror -async -syncgroup -name name | -groupid id
[-synctype startnow]
[-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group you want to update or synchronize. You can specify either the consistency group name or id (next).

`-groupid id`

Specifies the unique ID of the consistency group you want to update or synchronize.

`-synctype startnow`

Starts synchronization immediately. `startnow` is the only currently supported option.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command starts synchronizing the consistency group immediately.

```
naviseccli -h ssl_spa mirror -async -syncgroup -name dbgroup
-synctype startnow
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -fracturegroup

Fractures all image(s) in a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror -async` command with `-fracturegroup` starts an administrative fracture of all images in the consistency group, even if they are already system fractured. The consistency group cannot be in a scrambled state, and all images must be primary. If the consistency group is in the process of updating, it stops.

SYNTAX

```
mirror -async -fracturegroup -name name | -groupid id
[-fracturetype abort]
[-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to fracture. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to fracture.

-fracturetype abort

Specifies that you want to abort an update if one is in progress. The secondary retains any extents that have been copied to the secondary, and therefore do not need to be resent on the next synchronization.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, with the consistency group named `dbgroup`, this command fractures all the images in the consistency group.

```
naviseccli -h ss1_spa mirror -async -fracturegroup
-name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -promotegroup

Promotes all secondary images in a consistency group to primary images.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command promotes all secondary images to the role of the primary images. This command is used with `naviseccli` (see `naviseccli`).

SYNTAX

```
mirror -async -promotegroup -name name | -groupid id
[-type type] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to promote. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to promote.

-type *type*

Specifies the type of promotion. Values are:

`normal` (default), which causes a failure if the connection to the primary system is not functioning or if at least one set of the images to promote are not synchronized with the primaries. Otherwise, it will destroy the original mirrors and create new ones

with the targeted secondaries as primaries and the original primaries as secondaries.

local, which does not attempt to add back the original primaries as secondaries, but always creates new mirrors with only primary images consisting of the targeted secondaries.

-o

Runs the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command promotes all secondary images:

```
naviseccli -h ssl_spa mirror -async -promotegroup -name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -listgroups

Displays information about existing consistency groups.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror -async` command with `-listgroups` and no other switches lists all information for all asynchronous consistency groups on the system. Issuing this command with one or more options displays the information based on the option(s) issued. This command is used with `naviseccli` (see `naviseccli`).

SYNTAX

```
mirror -async -listgroups [-name[name]] [-groupid[-id]]
[-feature] [-update] [-nextupdate]
[-syncrate] [-state] [-role] [-mirrors]
[-condition][-recoverypolicy] [-description]
```

OPTIONS

`-name name`

Specifies the name of the consistency group whose information you want to display. If you do not specify the name of the consistency group with this switch, that piece of information will be provided for the consistency group(s) being displayed.

`-groupid id`

Specifies the unique ID of the consistency group whose information you want to display.

`-feature`

Displays the maximum number of consistency groups allowed on the system and the maximum number of members per consistency group.

`-update`

Displays the type of update as either manual, Time elapsed since the start of the last update (in minutes), or Time elapsed since the end of the last update (in minutes). On the next line it also displays the Update Period in minutes. If the Update Period is not applicable, it displays N/A.

`-nextupdate`

Displays the number of minutes until the next update.

`-syncrate`

- Displays the rate at which data is transferred during updates and initial synchronizations.
- state
 - Displays Synchronized, Out-of-Sync, Synchronizing, Consistent, Rollingback, Quasi-Consistent, Scrambled, Incomplete, Local Only, or Empty.
- role
 - Displays Primary, Secondary, or Unknown.
- mirrors
 - Displays the mirror name, mirror uid, the IDs of both images in the pair, and the system ID of the paired image.
- condition
 - Displays the consistency group condition as Normal, Initializing, Updating, Admin Fractured, System Fractured, Waiting on Admin.
- recoverypolicy
 - Displays either manual or automatic to describe whether the system will wait for you to restart updates in the event of a system fracture, after the problem is corrected.
- description
 - Displays the description associated with the consistency group.

EXAMPLE

```
naviseccli -h 10.14.5.200 mirror -async -listgroups
```

```
Group Name: taraFARGroup
Group ID: 50:06:01:60:90:20:5A:83:00:00:00:00
Description:
Maximum Number of Groups Allowed: 16
Maximum Number of Mirrors per Group: 16
Update Type: Manual
Time to next update: 0
State: Quasi Consistent
Role: Secondary
Mirror Name: taraFAR202
Mirror WWN: C6:05:20:88:60:01:06:50:08:00:00:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Mirror Name: taraFAR201
Mirror WWN: C6:05:20:88:60:01:06:50:07:00:00:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Mirror Name: taraFAR200
Mirror WWN: C6:05:20:88:60:01:06:50:06:00:00:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Condition: Admin Fractured - Intervention required
Sync Rate: Medium
Recovery Policy: Manual
Group Name: HWLocalPromoteTestGroup
Group ID: 50:06:01:60:90:20:5A:83:01:00:00:00
Description:
Maximum Number of Groups Allowed: 16
Maximum Number of Mirrors per Group: 16
Update Type: Manual
Time to next update: 0
State: Synchronized
Role: Secondary
Mirror Name: HWLocalPromoteTestMirror
Mirror WWN: C6:05:20:88:60:01:06:50:09:00:00:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Condition: Normal
```



```
Sync Rate: Medium  
Recovery Policy: Manual
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

CHAPTER 10

EMC VNX MirrorView/S CLI

This chapter introduces the EMC® VNX MirrorView™/Synchronous software and explains the commands for creating and managing the MirrorView/Synchronous remote mirrors and consistency groups. This manual refers to the EMC VNX MirrorView/Synchronous product as MirrorView/S.

Major topics include:

- [MirrorView/S Remote Mirroring commands](#)..... 444
- [MirrorView/S consistency group commands](#)..... 471

MirrorView/S Remote Mirroring commands

This section explains the `naviseccli` commands and the commands for creating and managing remote mirrors on VNX series systems with the MirrorView/S option. These commands let you use MirrorView/S software to create a byte-for-byte copy of one or more local LUNs connected to a distant system server.

Note

The commands in this section function only with a system that has the optional MirrorView/S software installed.

MirrorView/S operations overview

This section provides an overview of MirrorView/S operations and how to manage the systems using the management software.

Procedure

1. Connect the same management software to both systems and configure the software, so that you can manage both systems.

Note

Optionally, you can perform this step at the secondary site as well.

You must manage both systems, which can be in different domains (see the Unisphere online help).

2. If you want to use the write intent log, allocate the LUNs on all participating systems, so that if you promote, the write intent log is enabled.

In the CLI, to designate the write intent log, use the `mirror -sync -allocatelog` command. You can specify any LUN that can be made a private LUN, that is, any LUN that is not part of a storage group and is not a hot spare.

3. Establish a usable, two-way connection between the MirrorView/S systems using the `mirror -enablepath` command (see [mirror -enablepath on page 458](#)).
4. If the primary LUN does not exist, bind it on its server's system. Wait for the LUN to finish binding and add it to the storage group.
5. If the secondary LUN does not exist, create a secondary image LUN.

Note

The secondary LUN can be a different RAID type from the primary, as long as the block size matches.

6. Wait for the secondary LUN to finish binding.
7. On the host use the CLI to create the mirror using the `mirror -sync -create` function. You can specify write intent log information in the command. The `-create` function activates the mirror.
8. Create a consistency group and add the primary image to it (optional). See [mirror -sync -creategroup on page 471](#).
9. With the CLI, add the secondary images you need with the `-addimage` function. You can list images with the `mirror -sync -list` function.

By default, when you use the `-addimage` command, the software starts synchronizing the secondary image with the primary. Unless the source LUN is freshly bound or similarly contains no meaningful data, a full synchronization is necessary.

If you do not want the default action when you add the image (perhaps because the source LUN is newly bound and synchronizing the mirror to it would be pointless), then you can tell the CLI that synchronization is not required – an option in the `-addimage` command.

After you add the image, all writes to the primary LUN will be duplicated to the secondary LUNs.

At any time in the previous sequence of steps, you can get remote mirror status with the CLI `mirror -sync -info` or `-list` command. You can also check the progress of synchronization with the `-listsyncprogress` command.

10. If a primary failure occurs, Unisphere reports the failure.

If the primary failure is minor, have the primary fixed and resume mirroring.

If the primary failure is catastrophic, the original client that is managing the system may be unusable and thus unable to report the failure. For such a failure, the administrator at the secondary site must set up a client to manage the system (if not already done), then promote the secondary to primary and take other recovery action as needed. This includes assigning the newly promoted LUN to a storage group, if it is not already the appropriate one.

When you use MirrorView/S on a VMware ESX Server, after you promote the secondary image to a primary image, perform the following steps:

- a. If not already assigned, assign the newly promoted primary image to a storage group on the same or standby ESX Server.
- b. Rescan the bus at the ESX Server level.
- c. If not already created, create a virtual machine (VM) on the same or standby ESX Server. The VM is not powered up when you create it.
- d. Assign the newly promoted primary to the VM.
- e. Power up the VM.

If the VM is created and running and you have not already assigned the newly promoted primary to the VM, perform these steps:

- a. Perform steps a and b above.
- b. If you are running ESX Server 2.5.x, power it down.
- c. To assign the newly promoted primary to the VM, use the virtual center interface for ESX Server 3.x and 2.5.x or the Management User interface for ESX Server 2.5.x.

Note

If you are running ESX Server 3.x, you do not need to power down the VM.

- d. If you are running ESX Server 2.5.x, power up the VM.

Note

If you are running ESX Server 3.x, you do not need to power up the VM.

The primary image (which is now the secondary image) will not be accessible to the primary ESX Server.

Note

For configuration restrictions when using VMFS volumes, go to the E-Lab™ Navigator on EMC Online Support, and under the **PDFs and Guides** tab, open the `VMware ESX server.pdf` file.

11. If access to the secondary storage image fails, the primary system will fracture the remote mirror. If the problem with the secondary is minor (for example, replacing a cable), then the administrator can fix it. Mirroring will recover and resynchronize the image, if the image recovery policy is Automatic.

Whenever you want to stop mirroring, you can first fracture and remove the secondary images, and then destroy the mirror. This does not affect any data on either image LUN, and access to the primary LUN is also unaffected. The LUN that held the secondary image now becomes accessible as a regular LUN.

MirrorView/S CLI functions

The VNX CLI for Block provides one of the two interfaces to the MirrorView/S mirroring software. The Unisphere UI provides the other. You can use the `naviseccli mirror` command to set up and manage remote mirror sessions.

The mirror command functions are summarized in [naviseccli MirrorView/S command functions on page 446](#).

Table 44 naviseccli MirrorView/S command functions

Essential functions (in order performed)	
<code>mirror -enablepath</code>	Opens a path between the local and remote system. Only a single connection is required for synchronous and/or asynchronous mirroring.
<code>mirror -sync -create</code>	Creates a new remote mirror using an existing LUN. The LUN on which the <code>-create</code> command is executed becomes the primary image.
<code>mirror -sync -addimage</code>	Adds a secondary image to a mirror using a LUN on a remote system. Synchronization is optional — if specified, synchronization of the image begins automatically.
<code>mirror -sync -fractureimage</code>	Fractures a secondary image, suspending mirroring of data to that image. A fracture log is maintained, so that the image can be quickly resynchronized later.
<code>mirror -sync -removeimage</code>	Removes a secondary image from a mirror.
<code>mirror -sync -promoteimage</code>	Promotes a secondary image to primary, while simultaneously demoting the primary (if still accessible) to secondary. Required in order to access the data in the secondary image in the event of a failure of the primary.
<code>mirror -sync -syncimage</code>	Starts synchronization of the secondary image with the primary. Needed if automatic recovery is not on. I/O is allowed with the primary while the synchronization occurs. You can specify synchronization rate to avoid serious performance impact.

Table 44 navisecli MirrorView/S command functions (continued)

Essential functions (in order performed)	
<code>mirror -sync -destroy</code>	Terminates remote mirroring of a LUN. This does not affect the data in the LUN, and the host continues to access it uninterrupted.
Optional status functions (alphabetically) (see also <code>getall -mirrorview</code> in EMC Navisphere CLI Reference)	
<code>mirror -sync -info</code>	Displays MirrorView/S information about a system.
<code>mirror -sync -list</code>	Displays information about existing mirrors.
<code>mirror -sync -listlog</code>	Displays write intent log information.
<code>mirror -sync -listsyncprogress</code>	Displays the percentage of synchronization.
Optional reconfiguration functions (alphabetically)	
<code>mirror -sync -allocatelog</code>	Allocates LUNs as space for the write intent log. This is optional – it allows recovery through a partial resynchronization if a recoverable failure of the primary SP occurs.
<code>mirror -sync -change</code>	Changes remote mirror properties such as name and description.
<code>mirror -sync -changeimage</code>	Changes secondary image properties such as recovery policy and synchronization rate.
<code>mirror -sync -deallocatelog</code>	Removes the write intent log LUN that was allocated with the <code>-allocatelog</code> function. Allowed only when the log is not active.
<code>mirror -disablepath</code>	Severs all paths between the local system SPs and the remote system SPs.
<code>mirror -setfeature</code>	Enables or disables mirroring to/from a LUN. Required in only a few specific cases where the software has insufficient information to do this automatically.

mirror -sync -addimage

Adds an image to a previously created mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Note

When you add a secondary image, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to add a secondary image to a mirror using a LUN on a remote system. The `-addimage` function is required if you want mirroring to occur. You must direct the command to the system that holds the primary image. This command lets you create primary and secondary images on thin LUNs, if all the participating systems support mirroring on thin LUNs.

Note

The Thin Provisioning™ enabler must be installed to create thin LUNs.

You can use the `-info` function to display the UIDs of systems that have images that can be mirrored. Also, if the two arrays are not in the same domain, you need to use an IP address instead of a host name.

Note

To add a secondary image, you must have a working connection to the remote system, and the remote LUN must be exactly the same size (number of blocks) as the primary image.

Note

In this syntax "or" means that you can specify the switch, `-arrayhost sp-hostname | sp-IP-address -lun lun-number | -lunuid lun-uid` or you can specify the switch, `-arrayuid system-uid -lunuid lun-uid`. The rest of the syntax is common to both.

The options `-lun lun-number` and `-lunuid lun-uid` are optional. You can use either of the two options.

SYNTAX

```
mirror -sync -addimage -name name|-mirroruid mirroruid
-arrayhost sp-hostname -lun lun-number -lunuid lun-uid
[-recoverypolicy auto|manual]
[-syncrate high|medium|low] [-nosyncrequired]
```

or

```
mirror -sync -addimage -name name|-mirroruid mirroruid
-arrayuid system-uid -lunuid lun-uid
[-recoverypolicy auto|manual]
[-syncrate high|medium|low] [-nosyncrequired]
```

OPTIONS

`-name name`

Specifies the name that you gave to the mirror with the `-create` command. Use either this name or the mirror unique ID (16-byte world wide name).

`-mirroruid uid`

Specifies the unique 16-byte primary image ID (World Wide Name) the software assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the [mirror -sync -list on page 425](#) function).

`-arrayhost sp-hostname`

Specifies the hostname or network address of the remote SP that will hold the secondary image. The secondary image will be added to the remote mirror connected to this host's SP. Use this switch or `-arrayuid`.

`-lun lun-number`

Specifies the LUN number that will serve as the secondary image of the mirror.

`-arrayuid system-uid`

Specifies the unique 8-byte system ID of the remote SP. You can determine this by using the `mirror -info` command or the `getarrayuid` command. For any letters in the ID, use the same case that the software assigned when the mirror was created. Use this switch or `-arrayhost`.

`-lunuid lun-uid`

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to be added as a secondary image. If you specified `-arrayuid`, then `-lunuid` is required. If you specified `-arrayhost`, then you can use either `-lun` or `-lunuid`.

Note

When you specify `-arrayuid` and `-lunuid`, you must add the mirror driver on the remote LUN using the `-setfeature` command before executing the `-addimage` command.

`-recoverypolicy manual|auto`

Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:

`manual`: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

`auto`: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.

`-syncrate high|medium|low`

Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other system operations.

`-nosyncrequired`

Specifies that synchronization is not needed for the newly added secondary image. The image will not be useful until it is synchronized. If you omit this switch, a synchronization occurs.

EXAMPLE

Any of the following commands adds an image to a mirror. For SP `ss1_spa`, any of the above commands adds to the remote mirror `db_mirror` the LUN in system `50:06:06:10:00:FD:A1:69` (or a LUN on the remote system `cpc426`) whose unique LUN ID (world wide name, WWN) is `60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11` or whose LUN number is 9. See also the example in [mirror-sync - create on page 453](#).

```
naviseccli -h ss1_spa mirror -sync -addimage
-name db_mirror -arrayuid 50:06:06:10:00:FD:A1:69
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

```
naviseccli -h ss1_spa mirror -sync -addimage
-name db_mirror -arrayhost cpc426 -lun 9
```

```
naviseccli -h ss1_spa mirror -sync -addimage
-name db_mirror -arrayhost cpc426
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -allocatelog

Allocates or re-allocates a LUN for use as a write intent log.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror` command with `-allocatelog` assigns a LUN as the write intent log on each SP of the system. The recommended (and minimum) size for the log is 128 MB. The write intent log allows recovery with a partial resynchronization if a recoverable failure of the primary SP occurs. If you do not use the intent log, a full synchronization is always required in the event of a failure of the SP controlling the primary image.

You must specify both SPs in one command line.

SYNTAX

```
mirror -sync -allocatelog -spA LUN-number -spB LUN-number
[-unbind] [-o]
```

OPTIONS

`-spA LUN-number` and `-spB LUN-number`

Identify the LUN to be used for the write intent log. It must hold at least 128 MB, must not be part of a storage group and must not be a hot spare.

`-unbind`

If write intent log LUNs are currently defined, unbinds the old LUNs after the write intent log function starts using the new LUNs.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command allocates as the write intent log for SP A the LUN with ID 122 and for SP B the LUN with ID 124.

```
naviseccli -h ss1_spa mirror -sync -allocatelog -spA 122 -spB 124
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -change

Changes mirror properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to change one or more parameters of a previously created mirror.

You must direct the command to the system that holds the primary image. You must specify at least one of the optional switches with this command. If you omit the override option, then the CLI displays a message for each parameter to confirm the change.

Note

Always view and modify remote mirror properties from the primary system. Information displayed from the secondary system may not be up to date, especially if the primary system has lost contact with the secondary system.

Note

When you enable the write intent log on a mirror that did not previously use it, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -sync -change -name name | -mirroruid uid
[-newname name]
[-description description]
[-requiredimages num_of_images] [-o]
[-qthresh qthresh] [-usewriteintentlog yes|no]
```

OPTIONS

`-name name`

Specifies the name the `-create` command gave the mirror. You can specify either the mirror name or unique ID (next).

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

`-newname name`

Specifies the new name for the mirror. You can specify up to 32 characters.

`-description description`

Provides a meaningful title for the mirror (for example, Employee and Benefits Table for Human Resources). The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.

`-requiredimages num_of_images`

Specifies the minimum number of secondary images that should be accessible to the mirror. The default value is 0. Allowed values are 0, 1, and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.

`-qthresh qthresh`

Specifies the time delay, in seconds, after the last I/O is received before a mirror will automatically be transitioned to the synchronized state. The default is 60 seconds. Valid values are 10-3600.

`-usewriteintentlog yes|no`

Directs the software to use or not use the remote mirror write intent log. If you want to use the log, you must have allocated it with `-allocatelog`.

- yes
Use or start using, the write intent log.
- no
Do not use, or stop using, the write intent log.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For the remote mirror `dbmirror` on SP `ss0_spa`, this command reduces the minimum number of accessible secondary images to 0. Use the `-list` command to confirm that the changes were made.

```
naviseccli -h ss0_spa mirror -sync -change -name dbmirror
-requredimages 0
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -changeimage

Changes secondary image properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to change one or more parameters of a mirror image.

You must direct the command to the system that holds the primary image. You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

```
Do you want to change the sync rate from Medium to High (y/n)?
To change a parameter, enter y.
```

SYNTAX

```
mirror -sync -changeimage -name name | -mirroruid uid
-imageuid image-uid
[-recoverypolicy manual|auto]
[-syncrate high|medium|low] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the remote mirror in the `-create` command. You can specify either the mirror name or uid (next).

`-mirroruid uid`

Specifies the unique 16-byte ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

`-imageuid image-uid`

Specifies the unique 8-byte ID of the secondary image you want to change. For any letters in the ID, you must specify the same case that the software assigned to the image when the mirror was created (use the `mirror -list` command).

`-recoverypolicy manual|auto`

Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:

manual: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

auto: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.

`-syncrate high|medium|low`

Specifies a relative value (low, medium, or high) for the priority of completing updates. The default is medium. High completes updates faster, but may significantly affect system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other system operations.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, mirror `db_mirror`, image `50:06:06:10:00:FD:A1:6`, this command changes the synchronization rate to high.

```
naviseccli -h ssl_spa mirror -sync -changeimage -name db_mirror
-imageuid 50:06:06:10:00:FD:A1:6 -syncrate high
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -create

Creates a new mirror using an available LUN for mirroring.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Since the write intent log is on by default, before creating a synchronous mirror, you must first allocate the write intent log (see [mirror -sync -allocatelog on page 450](#)).

DESCRIPTION

This command is used with `naviseccli` to create a new mirror using an existing LUN. The LUN on which you run the command becomes the primary LUN. But, for mirroring to occur, you must add an image with the `-addimage` function. Later, you can change some of the parameters you specify here with the `-change` function. This command supports thin LUNs after you install the latest Unisphere version, which contains thin LUN support for MirrorView.

If you specify a minimum number of required images (other than the default value zero) in the `-create` command, the mirror will start in the attention state, since you have not yet added the required secondary images. To avoid this problem, do not specify a minimum number of images with `-create`; instead, create the mirror, add the image, and then if you want a minimum number of required images, change the mirror with the `-change` command.

You can use the `-info` command to display the LUN numbers that can be mirrored.

Note

When you create a mirror, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -sync -create -name name
-lun lun_number
[-description description]
[-requiredimages num_of_images]
[-qthresh qthresh]
[-nowriteintentlog ]
[-o]
```

OPTIONS

-name *name*

Lets you assign a meaningful name to the mirror; for example "Employee Benefits". To include spaces in the name, enclose the entire string in quotation marks as in the example. You can use this name to manage the mirror. The name can include as many as 32 characters.

-lun *lun_number*

Specifies the LUN number that will serve as the primary image of the mirror you are creating.

-description *description*

Lets you assign a description for the mirror. The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.

-requiredimages *num_of_images*

Specifies the minimum number of secondary images that should be accessible to this mirror. The default value is 0. Allowed values are 0, 1, 2, and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.

-qthresh *qthresh*

This specifies the time delay, in seconds, after the last I/O is received before a mirror will automatically transition to the synchronized state. The default is 60 seconds. Valid values are 10-3600.

-nowriteintentlog

Lets you create a synchronous mirror with the write intent log off. When creating a synchronous remote mirror, the write intent log will be on for the mirror by default.

-o

Executes the command without prompting for confirmation.

EXAMPLE #1

The mirror `-create` command creates the mirror named `db_mirror` on LUN 112.

```
naviseccli -h ss1_spa mirror -sync -create -name db_mirror -lun 112
```

EXAMPLE #2

The `-change` command sets a minimum of one secondary image for the mirror.

```
naviseccli -h ss1_spa mirror -sync -change -name db_mirror
-requiredimages 1
```

EXAMPLE #3

The `-list` command confirms the change.

```
naviseccli -h ssl_spa mirror -sync -list
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -deallocatelog

Removes LUNs from service as write intent log LUNs.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror` command with `-deallocatelog` deassigns the LUNs that have been used for the write intent log on this system. This command deallocates the LUNs used by both SPs of the system. The LUNs themselves and their contents are not affected by this command. If at least one mirror is still configured to use the write intent log LUNs, this command will fail with an error message.

If you omit the override option, it displays a message to confirm the deallocation of the intent log:

```
Do you want to deallocate write intent log (y/n)?
```

To deallocate the log for both SPs, answer **y**; to take no action, enter **n**.

SYNTAX

```
mirror -sync -deallocatelog [-unbind] [-o]
```

OPTIONS

`-unbind`

Directs the software to unbind the LUNs used in write intent logs. (optional)

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This command deallocates the LUNs used for the write intent log function on both `ssl_spa` and its peer. This command does not affect the LUNs themselves.

```
naviseccli -h ssl_spa mirror -sync -deallocatelog
```

```
Do you want to deallocate write intent log on SP A (y/n)?
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -destroy

Destroys a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to destroy a mirror. It is the counterpart of `-create`. It ends mirroring of the LUN that is part of the mirror. This does not affect the data in the LUN, and the server continues to access it.

Note

You cannot destroy a mirror while it is a member of a consistency group.

You must remove the secondary image in the mirror before issuing the command to the system holding the primary image. If you cannot remove the secondary image or access the primary image, you can use the `-force` option. You must direct the command to the system that holds the primary image unless `-force` is specified, in which case, you can direct it to either system.

CAUTION

Use `-force` only if all other options to remove an image and the mirror have failed.

If you omit the override option, a message to confirm the destruction of the mirror appears:

```
Do you really want to destroy the mirror <mirror-name> (y/n)?
To destroy the mirror, answer y; to take no action, enter n.
```

SYNTAX

```
mirror -sync -destroy -name name | -mirroruid uid
[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the mirror with the `-create` command. You can also use the remote mirror ID.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

`-force`

Forcibly destroys the mirror. This option affects only the local system.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, this command destroys the remote mirror that was created with the name `db_mirror`.

```
naviseccli -h ssl_spa mirror -sync -destroy db_mirror
```

```
Do you really want to destroy the mirror db_mirror?(y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages if the command fails for some reason.

mirror -disablepath

Disables the communication paths created by `-enablepath` to remote mirror SPs.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 20](#)) to sever the path established by `-enablepath` in the remote system.

When issued to any SP, the command severs paths for both SPs in both the primary and secondary systems; that is, it severs these paths if:

- SP A of the local system is connected to SP A of the remote system
- SP A of the remote system is connected to SP A of the local system
- SP B of the local system is connected to SP B of the remote system
- SP B of the remote system is connected to SP B of the local system

This means you can issue the command to either SP of either system to sever the paths between the systems you are targeting.

If you issue the command while a mirror exists, the CLI issues a warning message; if you continue, any mirror connected over the specified path will be fractured.

Note

If both asynchronous and synchronous mirrors are installed, you need to run the command only once. It will take effect for both mirrors.

If you omit the override option, the CLI displays a message to confirm the disabling of the path. To disable the path, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -disablepath connected-SP-hostname [-connectiontype
fibre|iscsi] [-o]
```

OPTIONS

connected-SP-hostname

Specifies the hostname of SP A or SP B of a connected system to disconnect.

`-connectiontype fibre | iscsi`

Specifies the connection type as fibre or iSCSI for the path you are disabling. If you do not specify the `-connectiontype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

`-o`

Executes the command without prompting for confirmation. Without this switch, the CLI displays a confirmation message for each SP.

EXAMPLE #1

For SP `ss1_spa`, this command and confirmation remove the paths between the local SPs and the remote mirror SPs.

```
naviseccli -h ss1_spa mirror -disablepath ss27_spa -connectiontype
fibre
```

Disabling of the path will stop mirroring of data to all images if any among these arrays. A system fracture of

```
all mirrors between two arrays on following storage
processors will occur due to this action.
128.221.39.29
10.14.20.154
Do you want to continue (y/n)? y
```

EXAMPLE #2

The `mirror -info` command displays the updated status.

```
naviseccli -h ssl_spa mirror -sync -info
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

mirror -enablepath

Establishes a path for remote mirroring between a local and remote system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 20](#)) to create a path between the specified SP and its peer SP to a remote SP and its peer. These paths will serve as paths for remote mirroring. A remote mirror requires a path between the primary and secondary systems.

When issued to any SP, the command creates paths for both SPs in both systems; that is, it creates these paths:

- SP A of the local system to SP A of the other system
- SP A of the other system to SP A of the local system
- SP B of the local system to SP B of the other system
- SP B of the other system to SP B of the local system

The SP As or the SP Bs in both systems must be physically connected.

If you cannot enable paths on one SP due to some error, an informational message appears and the paths are enabled for the other SP. To obtain status information, run the `mirror -info -systems` command.

Note

If both asynchronous and synchronous mirrors are installed, you will need to run the command only once. It will take effect for both mirrors.

Note

If the two arrays are not in the same domain, you need to use an IP address instead of a host name for remote mirroring.

SYNTAX

```
mirror -enablepath SP-hostname [-connectiontype fibre|iscsi]
```

OPTIONS

SP-hostname

Specifies the hostname of SP A or SP B of the remote system on which you want to enable remote mirroring.

```
-connectiontype fibre | iscsi
```

Specifies the connection type as fibre or iSCSI for the path you are enabling. If you do not specify the `-connectiontype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

Note

To change the connection type, you must first disable the current enabled path. To modify, test, or delete connections or change the credentials, see the *EMC Navisphere Command Line Interface (CLI) Reference*.

EXAMPLE

For SP `ss1_spa`, this command enables the paths between `ss1_spa` and its peer SP and `ss27_spa` and its peer SP. The `mirror -sync-info` command that follows displays the updated status.

```
naviseccli -h ss1_spa mirror -enablepath ss27_spa -connectiontype
iscsi
```

```
naviseccli -h ss1_spa mirror -sync -info
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

mirror -sync -fractureimage

Fractures an image in a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to start an administrative fracture of a secondary image. It also prevents scheduled updates from starting in the future. You must direct the command to the system that holds the primary image.

(An automatic fracture, called a system fracture, occurs if a secondary image becomes inaccessible to the primary image. After a system fracture, the system periodically tries to communicate with the secondary image and — if it succeeds — starts synchronization based on the fracture log. If you issue the `-fractureimage` command while the mirror is system fractured, the system will not try to communicate or resynchronize until you issue a `-syncimage` command.)

The software displays a message to confirm the fracture:

```
Warning! If the Write Intent Log is not enabled for this mirror, a
full mirror sync of the secondary image will be required in the event
of a failure on the primary image.
Do you want to fracture the image image UID (y/n)?
```

To fracture the mirror, answer **y**; to take no action, enter **n**.

SYNTAX

```
mirror -sync -fractureimage -name name | -mirroruid uid
-imageuid uid
[-o]
```

OPTIONS**-name** *name*Specifies the name you gave to the mirror with the `-create` command.**-mirroruid** *uid*Specifies the unique 16-byte mirror ID (World Wide Name) that the software assigned when it created the mirror. For any letters in the ID, see comments under `-imageuid` below.**-imageuid** *uid*Specifies the unique 8-byte image ID of the secondary image to be fractured. For any letters in the ID, you must specify the same case as the software assigned to the image when the mirror was created (use `mirror -list` function).**-o**

Executes the command without prompting for confirmation.

EXAMPLEFor SP `ssl_spa` of the mirror named `db_mirror`, this command fractures the secondary image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -sync -fractureimage -name
db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

```
Warning! If the Write Intent Log is not enabled for this
mirror, a full mirror sync of the secondary image will be
required in the event of a failure on the primary image.
Do you really want to fracture the image(y/n)? y
```

OUTPUTIf the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.**mirror -sync -info**

Displays general information about synchronous mirroring in a system.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTIONThis command, when used with `naviseccli` and no other switches, lists all general information for remote mirror support. Issuing this command with one or more options displays the information based on the option issued. See the example for sample output.For information on existing mirrors, use the `-list` function. For synchronization status information, use the `-listsyncprogress` function.**SYNTAX**

```
mirror -sync -info [-cancreate]
[-maxmirrors]
[-mirroredluns]
[-mirrorableluns]
[-systems [-connectiontype]]
[-logs]
[-all]
```

OPTIONS**-cancreate**

- Indicates whether an remote mirror can be created in this system: yes or no.
- maxmirrors
Displays the maximum number of mirrors that can be created in this system.
- mirroredluns
Lists the LUN numbers of LUNs that are mirrored in this system. The list includes LUNs used as both secondary and primary images.
- mirrorableluns
Lists LUNs that can be mirrored in this system but are not mirrored at present.
- systems
Lists remote systems that can be used for a secondary image.
- connectiontype
Displays the connection type (fibre or iSCSI) for the enabled mirror. For unexpected error cases, it displays the value, unknown.
- logs
Lists LUNs that are assigned for the write intent log on this system.
- all
Lists all properties of a session without having to type in any other switch.

EXAMPLE #1

For SP `ss1_spa`, this command displays potential and current remote mirror information.

```
naviseccli -h ss1_spa mirror -sync -info
```

```
Can a mirror be created on this system: YES
Logical Units that are mirrored in Primary Images: 82 76
Logical Units that are mirrored in Secondary Images: 75 5
72
Logical Units that can be mirrored: 10 6 18 74 16 9 85 11
87 77 19 14 250
78 4 15 7 83 17 12 88 100 71 13 89 81
Is Write Intent Log Used: NO
Remote systems that can be enabled for mirroring:
Remote systems that are enabled for mirroring:
Array UID: 50:06:01:60:20:04:A1:6B
Status: Enabled on both SPs
Array UID: 50:06:01:60:90:20:33:54
Status: Enabled on both SPs
Array UID: 50:06:01:60:20:02:C1:A3
Status: Enabled on both SPs
Enabled Connection Type(s): iSCSI
Disabled Connection Type(s): Fibre
Maximum number of possible Mirrors: 50
```

EXAMPLE #2

The following example shows the output for the `mirror -async -info` command.

```
mirror -sync -info -all
```

```
Can a mirror be created on this system: YES
Logical Units that are mirrored in Primary Images:
Logical Units that are mirrored in Secondary Images: 3
Logical Units that can be mirrored: 8186 8191 8188 8190 8187 8189
Is Write Intent Log Used: NO
Remote systems that can be enabled for mirroring:
Array UID: 50:06:01:60:88:60:01:8C
Available Connection Type(s): Fibre
Remote systems that are enabled for mirroring:
Array UID: 50:06:01:60:88:60:07:4B
Status: Enabled on both SPs
Enabled Connection Type(s): Fibre
Disabled Connection Type(s):
Maximum number of possible Mirrors: 256
Logged-in status of Remote systems:
```

```

Array UID: 50:06:01:60:88:60:01:8C
Fibre Logged-in Status: SPA:Y; SPB:Y
Array UID: 50:06:01:60:88:60:07:4B
Fibre Logged-in Status: SPA:Y; SPB:Y
MirrorView port(s):
SP A:
I/O Module Slot: 0
Physical Port ID: 0
Logical Port ID: 0
Port Type: Fibre
SP B:
I/O Module Slot: 0
Physical Port ID: 0
Logical Port ID: 0
Port Type: Fibre

```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Command is not supported error message is printed to stderr.

mirror -sync -list

Displays information about existing mirrors.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli -sync mirror` command with `-list` and no other switches lists all information for all mirrors on the system. Issuing this command with one or more options displays the information based on the option(s) issued. You can direct the command to any system; however, only the system holding the primary image of a mirror will always have the latest information about the mirror. Status obtained from other systems may be stale; do not use it unless the primary system is unavailable. See the example for sample output.

For synchronization status information, use the `-listsyncprogress` function.

SYNTAX

```

mirror -sync -list [-name [name]] [-mirroruid] [uid]
[-description]
[-faulted]
[-images]
[-imagecount]
[-imagesize]
[-lun]
[-qthresh]
[-requiredimages]
[-state]
[-transition]
[-usewriteintentlog]

```

OPTIONS

`-name` *name*

Specifies the name you gave the mirror with the `-create` function. You can use the mirror name or unique ID to obtain information on a mirror.

`-mirroruid` *uid*

- Specifies the unique 16-byte remote mirror ID (World Wide Name) the software assigned when it created the mirror.
- description
 - Lists the meaningful title for the mirror, if any. The software does not use the description to identify the mirror.
- faulted
 - Lists whether the remote mirror is faulted: yes or no. If it is faulted, the mirror is either fractured or in the attention state.
- images
 - Displays image information, such as the role of the image (primary or secondary), image ID, WWN of the LUN containing the image, recovery policy, sync rate, image state, image condition, and image transitioning.
- imagecount
 - Lists the current number of images that compose the mirror. This count includes the primary image.
- imagesize
 - Lists the size of the image (capacity of the LUN containing the image) in user blocks.
- lun
 - Lists the LUN number of the local image.
- qthresh
 - Lists the time delay, in seconds, after the last I/O is received before a mirror will automatically transition to the synchronized state. The default is 60 seconds.
- requiredimages
 - Lists the minimum number of secondary images required for the mirror.
- state
 - Describes the operational state of the mirror. Values are: inactive, active, and attention.
- transition
 - Specifies whether the mirror is in a changing state.
- usewriteintentlog
 - Lists whether the mirror uses the write intent log. You cannot use the log unless you first allocate it.

EXAMPLE

```
naviseccli -h ssl_spa mirror -sync -list
```

```
MirrorView Name: pbc sync on 82
MirrorView Description:
MirrorView UID:
50:06:01:60:90:20:29:85:01:00:00:00:00:00:00:00
Logical Unit Numbers: 82
Remote Mirror Status: Mirrored
MirrorView State: Active
MirrorView Faulted: YES
MirrorView Transitioning: NO
Quiesce Threshold: 60
Minimum number of images required: 0
Image Size: 204800
Image Count: 2
Write Intent Log Used: NO
Images:
  Image UID: 50:06:01:60:90:20:29:85
  Is Image Primary: YES
  Logical Unit UID:
60:06:01:60:57:A0:0E:00:56:7D:92:1C:37:73:D8:11
  Image Condition: Primary Image
  Preferred SP: A
Image UID: 50:06:01:60:90:20:33:54
Is Image Primary: NO
Logical Unit UID:
```

```
60:06:01:60:B7:7C:0E:00:02:DB:F3:EF:36:73:D8:11
Image State: Consistent
Image Condition: System fractured
Recovery Policy: Automatic
Preferred SP: A
Synchronization Rate: Medium
Image Faulted: YES
Image Transitioning: NO
Synchronizing Progress (%): N/A
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -listsyncprogress

Displays the percentage of synchronization that has occurred.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror` command with `-listsyncprogress` and no other switches lists all information for all remote mirrors. Issuing this command with one or more options displays the information based on the option included.

If you omit arguments, the CLI displays information on all mirrors.

You must direct the command to the system that holds the primary image; if not, the progress will display as `N/A`. Similarly, if an image is synchronized, the progress will display as `100`. See the example for sample output.

SYNTAX

```
mirror -sync -listsyncprogress [-name name | -mirroruid uid]
```

OPTIONS

`-name name`

Specifies the name given to the mirror with the `-create` command. You can use the mirror name or unique ID to obtain information on a particular mirror.

`-mirroruid uid`

Specifies the 16-byte mirror unique ID (world wide name) assigned when the mirror was created. If you specify this switch, information on the particular mirror appears.

EXAMPLE

```
naviseccli -h ssl_spa mirror -sync -listsyncprogress
```

```
MirrorView Name: pbc sync on 82
Has Secondary Images: YES
Image UID: 50:06:01:60:90:20:33:54
Image State: Consistent
Synchronizing Progress(%): N/A
```

```
MirrorView Name: pbc sync on LUN 75
Has Secondary Images: YES
Image UID: 50:06:01:60:90:20:29:85
Image State: Out-of-Sync
Synchronizing Progress(%): N/A
```


OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -listlog

Displays write intent log information.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror` command with `-listlog` lists the write intent log LUN numbers. If a write intent log is not allocated, then it displays nothing. You can direct this command to any system that has MirrorView/S installed.

Format of the display is:

```
SP:A/B
```

LUN-number (Number of LUN used as write-intent log)

SYNTAX

```
mirror -sync -listlog
```

EXAMPLE

For both `ss1_spa` and its peer SP, this command lists the write intent log LUN information. If the write intent log is allocated, example output may be as follows:

```
naviseccli -h ss1_spa mirror -sync -listlog
```

```
Storage Processor: SP A
LUN Number: 129
Storage Processor: SP B
LUN Number: 128
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` message is written to `stderr`.

mirror -sync -promoteimage

Promotes a secondary image to primary.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command promotes a secondary image to the role of the primary image. The command is successful if it can promote the secondary to a primary that is server-accessible. You must direct this command to a system that holds the secondary image — never to the system that holds the primary. The command will fail if you issue it to the system that holds the primary image.

You can promote a secondary image if the secondary is in a synchronized or consistent state. If the original primary is available, the software does a full resynchronization unless it confirms that the two images were synchronized.

Use `-promoteimage` carefully. The software will prompt for confirmation as follows before processing the command:

```
Warning! This command is a very powerful and unforgiving
command like formatting a disk. An improper use of this command
can create unexpected results. If you have not done any of the
following steps then answer 'n' and fulfil these steps and re-
issue the command. Did you put primary image in a quiescent
state? Is image synchronized? It does not have to be
synchronized but it is better to be synchronized to avoid any
kind of loss of data. Do you want to continue (y/n)?
```

Answering yes will continue with processing of promote command. Answering no stops the command execution.

When you promote a secondary image to primary, the software assigns a new mirror UID to the promoted image to distinguish it from the old mirror. The new status of the old primary image depends on whether the old primary image is accessible when promotion occurs:

- If the original primary image is not accessible when you issue the promote command, the software creates a new mirror with the original secondary image as the primary image, and the original primary is not part of the new mirror. If the original mirror had a second secondary, it is added as a secondary to the new mirror if it is accessible. If the original primary image becomes available later, you should destroy the original mirror using the Unisphere `force destroy`, or VNX CLI for Block `mirror -sync -destroy -force` option. After destroying the original mirror, you can (if desired) add the LUN that held the original primary image to the new mirror as its secondary image. This secondary image requires a full synchronization before it provides data protection.
- If the original primary is accessible when you issue the promote command, then the promoted image becomes primary and the original primary becomes secondary (that is, the images swap). The software then tests to see if the two images are *synchronized*. If it finds the images are *synchronized*, then it proceeds with mirrored I/O as usual. If the mirror is not in the *synchronized* state, then the mirror will be out-of-sync and waiting on admin. The recovery policy is changed to manual during the promotion, so that you must initiate the synchronization. If you choose to synchronize the mirror, it will be a full synchronization.

Note

A promotion changes the mirror's recovery policy to manual.

After promoting an image, you must explicitly add the newly promoted image to an appropriate storage group. EMC recommends that you also remove the old primary from its storage group, if possible.

Note

You cannot issue this command on a mirror that is part of a consistency group.

Note

Promoting when I/O is going to the primary image can cause data loss. Any I/Os in progress during the promotion may not be recorded to the secondary image and will be unavailable after the secondary image is promoted. It will also probably require a full synchronization of the new secondary image after the promotion.

In normal operation when there is no failure, you can (but probably would not want to) promote a secondary image. If you do so, the primary image is demoted to secondary.

The new primary image inherits the mirror properties of the demoted primary image. Before issuing this command consider the following:

- If possible, the secondary image should be synchronized (synchronized state) with the primary image. If the secondary image is in the consistent state (not synchronized state) when you promote it to primary image, data will probably be lost during the promotion process.
The secondary image to be promoted must be in a consistent or synchronized state. If the secondary image is out-of-sync or synchronizing, then you cannot promote it; an error will result.
- You must stop I/O to an accessible primary image before you issue this command. If I/O continues to the original primary, data may be lost as the software demotes the primary to secondary and/or an SP may reboot unexpectedly during the attempted promotion.
- If the primary system has a write intent log, ensure that the secondary system does also.

SYNTAX

```
mirror -sync -promoteimage -name name | -mirroruid uid
-imageuid imageuid
[-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-create` command. You can use the mirror name or unique ID to promote a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte mirror ID (World Wide Name) that was assigned when the mirror was created. For any letters in the ID, see comments under `-imageuid` below.

`-imageuid uid`

Specifies the unique 8-byte image ID of the secondary image you want to promote. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the `mirror -list` function).

`-o`

Runs the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the remote mirror named `db_mirror`, this command promotes the secondary image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -sync -promoteimage -name db_mirror
-imageuid 50:06:06:10:00:FD:A1:6
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported error message` is printed to `stdout`. You may also receive other error messages.

mirror -sync -removeimage

Removes a secondary image from a mirror.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to remove a secondary image from the remote mirror. You must issue this command to a system that holds the primary image.

An image can be removed even if it is an active image. However, you can remove a secondary image only if it is in the synchronized state or you have fractured the mirror.

If you omit the `-override` option, the software displays a message to confirm the remove operation:

```
Removing of the image will stop mirroring of data to the
specified image. It may cause the minimum number of images
required parameter to be violated that will cause the mirror to
go into attention state. Do you really want to remove the
image(y/n)?
```

Answering yes will remove the image.

Note

You cannot issue this command on a mirror that is part of a consistency group.

SYNTAX

```
mirror -sync -removeimage -name name | -mirroruid mirroruid
-arrayhost sp-hostname | -imageuid imageuid [-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-create` command. You can use the mirror name or unique ID to remove a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

`-arrayhost sp-hostname`

Specifies the hostname or network address of the remote SP that holds the secondary image you want to remove. If you specify this switch, the driver(s) from the secondary LUN stack is removed.

`-imageuid image-uid`

Specifies the unique 8-byte image ID of the secondary image you want to remove. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the `mirror -list` command). If you specify this switch, you must explicitly remove the driver(s) from the secondary LUN stack using the `-setfeature` command.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP address `ssl_spa`, the owner of the mirror named `db_mirror`, this command removes the image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -sync -removeimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

```
Removing of the image will stop mirroring of data to the
specified image. It may cause the minimum number of
images required parameter to be violated that will cause
```

```
the mirror to go into attention state. Do you really want
to remove the image (y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -setfeature

Adds or removes an synchronous mirror driver to or from the LUN stack.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see `naviseccli`) to add the synchronous mirror driver to or remove it from the LUN stack depending on whether you specify the `-on` or `-off` option. You must issue this command to the system where the LUN resides.

This command is required after you remove a secondary image from a mirror, using the `-imaguid` switch to specify the image. In this case, Unisphere has insufficient information to automatically remove the MirrorView/S driver from the stack of the secondary image LUN, and you must use this command to do this manually.

Note

When you add a synchronous mirror driver to the LUN stack using the `-on` switch, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -sync -setfeature -on | -off -lun lun_number | -lunuid
uid
```

OPTIONS

`-on`

Adds the driver to the LUN stack.

`-off`

Removes the driver from the LUN stack.

`-lun lun_number`

Specifies the LUN number to which you want to add or from which you need to remove the driver.

`-lunuid uid`

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to which you will add the driver or from which you will remove the driver.

EXAMPLE

For SP `ss1_spa`, this command starts adding the driver to LUN 112.

```
naviseccli -h ss1_spa mirror -sync -setfeature -on -lun 112
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -sync -syncimage

Starts an update of the secondary image.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

Use this command with `naviseccli` (see `naviseccli`) to begin the synchronization operation on a secondary image that was previously administratively fractured or on one where the mirror image condition is normal. You must issue this command to the system that holds the primary image. The secondary image should be fractured when performing this operation. If you specify a secondary image that is not fractured, an internal error message will be returned.

If `-o` is not specified, the software displays a message asking you to confirm:

During synchronization, array performance will be impacted and data is vulnerable to a failure of the primary image. Do you want to sync the image (y/n)?

Answering **y** will start synchronization.

Note

You cannot issue this command on a mirror that is part of a consistency group.

SYNTAX

```
mirror -sync -syncimage -name name|-mirroruid mirroruid
-imageuid imageuid [-o]
```

OPTIONS

`-name name`

Specifies the name given to the mirror with the `-create` command. You can use the mirror name or unique ID to update a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

`-imageuid imageuid`

Specifies the unique 8-byte image ID of the secondary image you will synchronize. For any letters in the ID, you must specify the same case as the software assigned to the image when the software created the mirror (use the `mirror -list` command).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the mirror named `db_mirror`, this command starts synchronizing the image with the unique ID shown.

```
naviseccli -h ssl_spa mirror -sync -syncimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

During synchronization, array performance will be impacted and data is vulnerable to a failure of the primary image. Do you want to sync the image (y/n)? y

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Command is not supported error message is printed to stderr. You may also receive other error messages.

MirrorView/S consistency group commands

This section explains the commands for creating and managing groups of mirrors on VNX series systems with the EMC MirrorView/S option. This section refers to groups of mirrors as *consistency groups*.

Note

The LUNs in a MirrorView/S consistency group must reside on the same system, unlike the volumes in a Symmetrix consistency group, which can reside on multiple systems.

mirror -sync -creategroup

Establishes a consistency group to which you can later add mirrors.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

This command is used with `naviseccli` to create a consistency group to which you can add mirrors with subsequent commands. You must direct the command to the system on which you are creating the MirrorView/S group.

Note

You should not enable the FAST Cache for write intent log LUNs and Clone Private LUNs. Enabling the FAST Cache for these LUNs is a suboptimal use of the FAST Cache and may reduce the performance of the cache for other LUNs.

When you create a consistency group, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -sync -creategroup -name name
[-description description] [-recoverypolicy policy] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group. Use only alphanumeric characters for the consistency group name. The name can be up to 32 characters.

`-description description`

Specifies a textual description of the consistency group.

`-recoverypolicy policy`

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

manual: Update of the secondary image does not begin until you explicitly issue a `synchronize` command.

auto: Update of the secondary image automatically begins as soon as the connection between the two systems is again operational.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command creates a consistency group with the name `dbgroup`.

```
naviseccli -h ss1_spa mirror -sync -creategroup -name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -destroygroup

Destroys a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to destroy a consistency group. It is the counterpart of `-creategroup`.

Note

Make sure that the consistency group is empty before destroying it.

If you omit the override option, the system displays a message to confirm the destruction of the consistency group:

```
Do you really want to destroy the group <group-name>? (y/n)?
```

To destroy the consistency group, answer **y**; to take no action, enter **n**.

SYNTAX

```
mirror -sync -destroygroup -name name | -groupid id
[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-creategroup` command. You can also use the consistency group ID.

`-groupid id`

Specifies the consistency group unique ID assigned when the consistency group was created.

`-force`

Destroys the consistency group even if it is not empty. Force destroy destroys only the group on the local system; you must also destroy the group on the other system.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, this command destroys the consistency group that was created with the name `dbgroup`.

```
naviseccli -h ssl_spa mirror -sync -destroygroup dbgroup
```

```
Do you really want to destroy the group dbgroup?(y/n)? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages if the command fails.

mirror -sync -addtogroup

Adds a mirror to a previously created consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to add one mirror at a time to a consistency group.

Note

The mirror you specify must have its primary image on the system. No two secondary images of mirrors in the consistency group can be on different systems. When you add a mirror to a consistency group, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -sync -addtogroup -name name | -groupid id  
-mirrorname name | -mirroruid id
```

OPTIONS

`-name name`

Specifies the name of the consistency group to which you want to add the mirror.

`-groupid id`

Specifies the unique ID of the consistency group to which you want to add the mirror.

`-mirrorname name`

Specifies the name of the mirror you want to add to the consistency group.

`-mirroruid id`

Specifies the World Wide Name of the mirror you want to add to the consistency group.

EXAMPLE

For SP `ssl_spa`, this command adds the mirror, `Sales Accounts Mirror to Europe`, to the consistency group, `dbgroup`.

```
naviseccli -h ssl_spa mirror -sync -addtogroup -name dbgroup  
-mirrorname "Sales Accounts Mirror to Europe"
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -removefromgroup

Removes a mirror from a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` (see `naviseccli`) to remove one mirror at a time from the consistency group.

Note

The mirror you specify must have its primary image on the system. This does not apply if you use the `-force` switch.

If you omit the `-o` option, the software displays a message to confirm the remove operation. Answering yes will remove the mirror.

SYNTAX

```
mirror -sync -removefromgroup -name name|-groupid id
-mirrorname name|-mirroruid id[-force] [-o]
```

or

```
mirror -sync -removefromgroup -name name|-groupid id
-localLunUid uid | -localLunnumber number[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group from which you want to remove the mirror.

`-groupid id`

Specifies the unique ID of the consistency group to which you want to add the mirror.

`-mirrorname name`

Specifies the name of the mirror you want to remove from the consistency group.

`-mirroruid id`

Specifies the World Wide Name of the mirror you want to remove from the consistency group.

`-localLunUid uid`

Specifies the world wide name of the local LUN belonging to the local image, which participates in the consistency group. You can use this switch to remove group members only in the failure scenario.

`-localLunnumber number`

Specifies the LUN number of the local LUN belonging to the local image, which participates in the consistency group. You can use this switch to remove group members only in the failure scenario.

Note

Failure scenarios may occur where a mirror may no longer exist, but is still considered to be a member of the consistency group. In this case, the consistency group is in an incomplete state. A mirror can be missing from the consistency group if the promotion fails between destroying the original mirror and creating a new one. In this case, all the properties of the missing mirrors, except the `local LUN uid` and `local LUN number` appear as unknown in the `-listgroups` command.

You can remove the mirror from the consistency group only by specifying the group name or group ID and the local LUN uid or local LUN number. Use the local LUN uid or local LUN number only in this failure case. Removing a regular mirror in a group by specifying `-local lun number` or `-local lun uid` returns an error. Removing the missing mirror by specifying the `-mirror name` or `-mirror uid` also returns an error.

If a mirror is missing only on the secondary side, then issuing the `removefromgroup` command on the primary side will remove the missing member on the secondary side also.

-force

Removes the mirror even if the connection with the remote system is not functioning. If you direct this command at the secondary system with this switch, only the secondary member is removed; the primary system is not contacted.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command removes the mirror with the unique ID shown.

```
naviseccli -h ssl_spa mirror -sync -removefromgroup
-name dbgroup -mirroruid
```

```
83:5A:20:90:60:01:06:50:01:00:00:00:00:00:00:00
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -change group

Changes consistency group properties.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

This command is used with `naviseccli` to change one or more parameters of a consistency group.

You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

```
Do you want to change the group synchronization rate from
Medium to High (y/n)?
```

To change a parameter, enter **y**.

SYNTAX

```
mirror -sync -change group -name name | -groupid id
[-newname newname]
[-description description]
[-recoverypolicy policy] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to change. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to change.

-description *description*

Specifies a textual description for the consistency group.

-recoverypolicy *policy*

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

manual: Update of the secondary image does not begin until you explicitly issue a `synchronize` command.

auto: Update of the secondary image automatically begins as soon as the connection between the two systems is again operational.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For system `ssl_spa`, consistency group `dbgroup`, this command automatically starts synchronization.

```
naviseccli -h ssl_spa mirror -sync -change group -name dbgroup
-recoverypolicy auto
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -syncgroup

Synchronizes all members of the consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror -sync` command with `-syncgroup` synchronizes all members of a consistency group. When one member of the group is fractured, all will be fractured; therefore, it is necessary to synchronize the group. You cannot synchronize an individual image whose primary image is part of a consistency group. The `-syncgroup` command returns successfully that all members of the consistency group have been queued for synchronization.

SYNTAX

```
mirror -sync -syncgroup -name name | -groupid id
[-o]
```

OPTIONS

- name *name*
Specifies the name of the consistency group you want to update or synchronize. You can specify either the consistency group name or id (next).
- groupid *id*
Specifies the unique ID of the consistency group you want to update or synchronize.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command starts synchronizing the consistency group.

```
naviseccli -h ssl_spa mirror -sync -syncgroup -name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -fracturegroup

Fractures all image(s) in a consistency group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror -sync` command with `-fracturegroup` starts an administrative fracture of all images in the consistency group, even if they are already system fractured. You must direct this command to the primary system.

SYNTAX

```
mirror -sync -fracturegroup -name name | -groupid id [-o]
```

OPTIONS

- name *name*
Specifies the name of the consistency group you want to fracture. You can specify either the consistency group name or ID (next).
- groupid *id*
Specifies the unique ID of the consistency group you want to fracture.
- o
Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, with the consistency group named `dbgroup`, this command fractures all the images in the consistency group.

```
naviseccli -h ssl_spa mirror -sync -fracturegroup -name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -promotegroup

Promotes all secondary images in a consistency group to primary images.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

Before promoting, perform the following steps:

1. Stop all I/O.
2. If possible, make sure that the group is in a synchronized state.
3. Before running a force promote, first perform a normal promotion and, if possible, fix any error conditions that it returns.

⚠ CAUTION

After force promoting a consistency group that is out-of-sync after the promotion, use caution before synchronizing the consistency group. If the group was fractured before the promotion, the new primary system may have data that is much older than the original primary (now secondary) system. Synchronizing the group may overwrite newer data on the original primary system.

DESCRIPTION

This command promotes all secondary images to the role of the primary images. You must direct this command to a secondary system. This command is used with `naviseccli` (see `naviseccli`).

Note

To promote a consistency group, the group on the secondary side must be Consistent or Synchronized. When the primary system is not reachable, the group may not be in a Consistent or Synchronized state.

⚠ CAUTION

Promoting when I/O is going to the primary image can cause data loss. Any I/Os in progress during the promotion may not be recorded to the secondary image and will be unavailable after the secondary image is promoted. It will also probably require a full synchronization of the new secondary image after the promotion.

SYNTAX

```
mirror -sync -promotegroup -name name | -groupid id
[-type type] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group you want to promote. You can specify either the consistency group name or ID (next).

`-groupid id`

Specifies the unique ID of the consistency group you want to promote.

`-type type`

Specifies the type of promotion. Values are:

`normal` (default), which promotes each mirror in the consistency group. Otherwise, it will destroy the original mirrors and create new ones with the targeted secondaries as primaries and the original primaries as secondaries.

local, which creates only the new primary image. If the original primary system can be contacted, then the secondary image is removed from the mirror on the original primary system. Thus, both systems will have a mirror with primary images only.

force, which continues the promotion despite most errors that may occur. We recommend that before using force promote, first use normal promote and, if possible, fix any errors that it returns.

-o

Runs the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, the owner of the consistency group named `dbgroup`, this command promotes all secondary images:

```
naviseccli -h ssl_spa mirror -sync -promotegroup -name dbgroup
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -listgroups

Displays information about existing consistency groups.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli mirror -sync` command with `-listgroups` and no other switches lists all information for all synchronous consistency groups on the system. Issuing this command with one or more options displays the information based on the option(s) issued. This command is used with `naviseccli` (see `naviseccli`).

SYNTAX

```
mirror -sync -listgroups [-name[name]] [-groupid[-id]]
[-feature] [-state] [-role] [-mirrors]
[-condition] [-recoverypolicy] [-canbeaddedtogroup]
```

OPTIONS

`-name` *name*

Specifies the name of the consistency group whose information you want to display. If you do not specify the name of the consistency group with this switch, that piece of information will be provided for the consistency group(s) being displayed.

`-groupid` *id*

Specifies the unique ID of the consistency group whose information you want to display.

`-feature`

Displays the maximum number of consistency groups allowed on the system and the maximum number of members per consistency group.

`-state`

Displays Synchronized, Out-of-Sync, Synchronizing, Consistent, Scrambled, Incomplete, Local Only, or Empty.

`-role`

Displays Primary, Secondary, or Unknown.

`-mirrors`

Displays the mirror name, mirror uid, the IDs of both images in the pair, and the system ID of the paired image.

- condition
Displays the consistency group condition as Active, Admin Fractured, System Fractured, Waiting on Sync, Inactive, or Invalid.
- recoverypolicy
Displays either manual or automatic to describe whether the system will wait for you to restart updates in the event of a system fracture, after the problem is corrected.
- description
Displays the description associated with the consistency group.
- canbeaddedtogroup
Displays all the mirrors that you can add to a group. A mirror appears in the list only if it is consistent or synchronized and not fractured. Also, the consistency group must be in-sync or consistent and not fractured.

EXAMPLE

```
naviseccli -h 10.14.5.200 mirror -sync -listgroups
```

```
Maximum Number of Groups Allowed: 16
Maximum Number of Mirrors per Group: 16

Group Name: vid
Group ID: 50:06:01:60:90:60:2D:03:00:00:00:00
Description:
State: Synchronized
Role: Secondary
Condition: Active
Recovery Policy: Manual
Mirror Name: vid
Mirror WWN:
50:06:01:60:90:60:2C:BA:01:00:00:00:00:00:00:00:00
Primary WWN: 50:06:01:60:90:60:2C:BA
Secondary WWN: 50:06:01:60:90:60:2D:03
Array hosting remote image: 50:06:01:60:90:60:2C:BA
Local LUN ID:
60:06:01:60:8B:53:10:00:A0:FF:D5:30:DE:B4:D8:11
Local LUN Number: 0
Remote LUN ID:
60:06:01:60:90:53:10:00:A3:58:04:80:DE:B4:D8:11
Mirrors can be added to group:
Mirror Name: djp1
Mirror WWN:
50:06:01:60:90:60:2C:BA:03:00:00:00:00:00:00:00:00
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

CHAPTER 11

SAN Copy and its command line utilities

This chapter introduces the EMC® SAN Copy™ software and its command line utilities: admhost, with commands that the operating system uses to manage SAN Copy devices; and the CLI, with commands that manage SAN Copy sessions.

Major topics are:

•	Introducing SAN Copy software.....	482
•	Preparing a Block system for a copy session.....	482
•	Preparing a Symmetrix system for a copy session.....	485
•	About admhost for SAN Copy.....	488
•	admhost SAN Copy command functions.....	488
•	Preparing logical units for a SAN Copy session.....	490
•	SAN Copy CLI commands.....	492
•	SAN Copy command functions.....	492
•	sancopy -create.....	494
•	sancopy -destinfo.....	498
•	sancopy -duplicate.....	501
•	sancopy -info.....	502
•	sancopy -mark.....	506
•	sancopy -modify.....	507
•	sancopy -pause.....	512
•	sancopy -remove.....	513
•	sancopy -resume.....	514
•	sancopy -settings.....	515
•	sancopy -start.....	516
•	sancopy -stop.....	518
•	sancopy -throttle.....	519
•	sancopy -transfer.....	520
•	sancopy -unmark.....	521
•	sancopy -updateconnections.....	521
•	sancopy -verify.....	522

Introducing SAN Copy software

EMC SAN Copy software (referred to as SAN Copy later in this document) runs on a SAN Copy system (system with SAN Copy software enabled). It copies data between block systems, within block systems, between block and Symmetrix systems, and between block and qualified non-EMC systems. SAN Copy can copy data directly from a source logical unit on one system to destination logical units on other systems, without using host resources. SAN Copy connects directly or through a SAN, and also supports protocols that let you use the IP WAN (wide area network) to send data over extended distances. SAN Copy can perform multiple copies—each in its own copy session—simultaneously. The RAID type of the logical units participating in a copy session does not have to be the same; that is, the source and destination logical units can be different RAID types.

You can use SAN Copy to create full and incremental copies of a source logical unit. An incremental session copies only the data that has changed since the last copy session. This can significantly reduce the time needed to copy the data, thereby allowing the copy operation to be performed more frequently and more effectively. Unlike full copy sessions, the source logical unit for an incremental session can remain online during the copy process. You must, however, prepare the source logical unit prior to starting the session (see [Preparing logical units for a SAN Copy session on page 490](#)). Regardless of the type of copy session—full or incremental—the participating logical units must meet certain requirements.

You can use SAN Copy for the following tasks:

- One-time data migration
- Routine copying of data for application testing
- Content distribution
- Scheduled updates to multiple databases

To manage SAN Copy sessions, you can use either the UI-based Unisphere (described in the Unisphere online help), or the SAN Copy command line interface described in this manual.

Preparing a Block system for a copy session

If you will be copying data between Block systems or between VNX Block and Symmetrix systems, follow the SAN Copy configuration requirements in *EMC VNX for Block Concepts*, and prepare the Block systems as described in the following procedure. You can use logical units (LUNs) in a Block system as the source or destination of SAN Copy sessions.

If you will be copying data between LUNs within a Block system, see "Copying data between LUNs within a Block system" described in *EMC VNX for Block Concepts*.

Note

For all sessions, you must install SAN Copy software on at least one of the Block systems. For incremental copy sessions, the source logical unit must reside in the SAN Copy system.

Procedure

1. For Fibre Channel replication, use the native switch management tools to zone at least one port from an SP on the SAN Copy system to one or more SP ports on the participating remote systems (see Fibre Channel zoning requirements and recommendations described in *EMC VNX for Block Concepts*).

For iSCSI replication, create connection sets. A connection set establishes a replication session between an iSCSI initiator and an iSCSI target. A connection set consists of a unique user-defined name, optional username and password (secret), method for assigning security credentials (connection specific, array-shared or none), and one or more connection paths. Each connection set should include only one iSCSI initiator.

Note

SAN Copy ports act as host initiators to remote systems (system without SAN Copy software).

2. Register the SAN Copy ports with any other SP ports in the same zone using the `naviseccli` or `navicli -h spname sancopy -updateconnections` command (see [sancopy -updateconnections on page 521](#)).

If the remote system is a Block system, the SAN Copy ports are registered with each system, and the software adds a record for each SAN Copy port to the system's host initiator table. Unisphere displays an entry for each registered SAN Copy port in the Connectivity Status dialog box, and the SAN Copy Summary dialog box.

3. Verify that the SAN Copy ports are connected and registered with any systems that will participate in a copy session using the following command for each SP on the SAN Copy system:

```
naviseccli -h spname -port -list
```

where *spname* is the name of an SP in the SAN Copy system.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

4. Connect the SAN Copy ports to Block groups (see the next section, [Connecting SAN Copy ports to storage groups on page 483](#)).

Connecting SAN Copy ports to storage groups

Note

Each SAN Copy port acts like a host initiator and, therefore, can connect to only one storage group at a time in a storage system.

In order for the LUNs to be accessible to the SAN Copy ports, you must connect the SAN Copy ports to the storage groups in which the LUNs reside. When connecting to the storage group, SAN Copy lets you select all available SP ports with valid connections to the storage group, or it lets you select specific ports.

If all participating logical units of a remote storage system already belong to the same storage group, you can connect SAN Copy ports to the existing storage group. However, if the logical units belong to different storage groups, we recommend that you create one or more new storage groups, place the logical units in these storage groups, and then connect SAN Copy ports to them. This is the only time we recommend placing the same logical unit in more than one storage group.

Note

Some operating systems do not support the same host accessing identical LUNs. Therefore, once SAN Copy completes the copy session, we recommend that you move any LUN copies to different storage groups.

For more information on the following CLI commands, refer to the *EMC Command Line Interface (CLI) Reference*. To use the `navicli` CLI command, replace `naviseccli` with `navicli`.

Procedure

1. Use the following command to create storage groups:

```
naviseccli -h sname storagegroup -create -gname
storagegroupname
```

2. Use the following command to add LUNs to storage groups:

```
naviseccli -h sname storagegroup -addhlu -gname
storagegroupname -hlu hlunnumber -alu alunnumber
```

3. Use the following command (with information from the previous `naviseccli -port -list` command) to add the SAN Copy ports to storage groups:

```
naviseccli -h sname storagegroup -setpath -gname
storagegroupname -hbawwn hbawwn -sp a or b -sport
portnumber
```

Table 45 naviseccli -h Parameters

Parameter	Description
<code>sname</code>	IP address of the SP on the storage system to which the storage group belongs.
<code>storagegroupname</code>	name of the storage group to which you want to connect.
<code>hbawwn</code>	World wide name of the SAN Copy port (from the <code>-portlist</code> command) .
<code>a or b</code>	SP that owns the storage group.
<code>portnumber</code>	the ID of the SP port hat owns the storage group (from the <code>-portlist</code> command).

If the sharing state of the storage group is Dedicated (only a single host can access the LUNs in the storage group), you must also specify one of the following switches:

```
-host hostname, or -ip ipaddress
```

where:

Table 46 host Switches

Parameter	Description
<code>hostname</code>	name of the host that has access to the LUNs in the storage group
<code>ipaddress</code>	IP address of the host that has access to the LUNs in the storage group

[Sample SAN Copy storage group configuration on page 485](#) shows a sample storage group configuration for SAN Copy. The remote storage system includes three storage

groups labeled SG 1, SG 2 and SG 3. SG 1 and SG 2 provide I/O access to the LUNs from the attached hosts. SG 3 is a SAN Copy storage group.

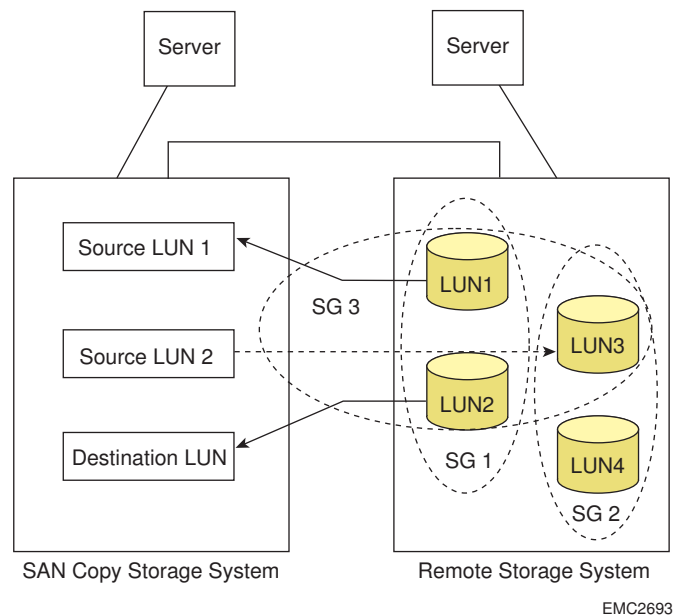
In this example, we start with two SAN Copy sessions running concurrently. The LUNs in storage group 1 (LUN 1 and LUN 2) are participating in both sessions. Source LUN 1 on the SAN Copy storage system is copying data to LUN 1 in SG 1, and LUN 2 in SG 1 is copying data to the destination LUN on the SAN Copy system. Since both LUN 1 and LUN 2 belong to the same storage group, you can connect the same SAN Copy port directly to this storage group and both sessions can use this SAN Copy port.

In the future, you may want to start a third session that copies data from source LUN 2 to LUN 3 using the same SAN Copy port. Since LUN 3 is in a different storage group, we recommend that you create a new storage group (SG 3) for the SAN Copy sessions, and place LUN 1, LUN 2 and LUN 3 in this storage group. This is the only time we recommend placing the same logical unit in more than one storage group.

Note

Some operating systems do not support the same host accessing identical LUNs. Therefore, once SAN Copy completes the copy session, we recommend that you move any LUN copies to different storage groups.

Figure 3 Sample SAN Copy storage group configuration



EMC2693

Preparing a Symmetrix system for a copy session

You can use logical units (volumes) in Symmetrix systems as the source or destination of a full SAN Copy session. For incremental sessions, the source logical units must reside in a SAN Copy system. In order to prepare a Symmetrix system for SAN Copy sessions, you must follow the SAN Copy configuration requirements section in *EMC VNX for Block Concepts*, and use one of the following procedures.

Preparing a Symmetrix system using EMC ControlCenter SAN Manager

Note

To complete this procedure, you must be familiar with Unisphere, EMC ControlCenter[®] SAN Manager, and Solutions Enabler Device Manager CLI.

Procedure

1. Use EMC ControlCenter SAN Manager 5.2 or higher to zone at least one port from an SP on the SAN Copy system to one or more FA ports on the participating Symmetrix systems (see Fibre Channel zoning requirements and recommendations section in *EMC VNX for Block Concepts*).
-

Note

You can also use the native switch management tools to zone a SP to one or more Symmetrix FA ports.

2. Verify that the host:
 - Has access to the Symmetrix volumes that are participating in the SAN Copy session
 - Is running the Unisphere Host Agent
 - Is connected to a system in the domain
 - Is part of a portal configuration within the domain. For information about portals, refer to the Unisphere online help.
3. Use Unisphere to register the SAN Copy ports with Symmetrix systems:
 - a. In the Storage tree of the Enterprise Storage window, right-click the icon for the SAN Copy system.
 - b. Select **SAN Copy > Update Connections**.
The SAN Copy ports log in to the Symmetrix storage ports.
4. To verify that the SAN Copy ports have been registered with the Symmetrix system, use the Solutions Enabler CLI command:

```
symmask -sid symmID list logins
```

where *symmID* is the Symmetrix serial ID.

Note

Use the SymmCLI command `symcpd list -sid` to determine the pathname for the VCM database, or run `inq` on the host, if available.

5. Use the following Solutions Enabler `symmask` command to assign Symmetrix volumes to the SAN Copy ports:

```
symmask -sid symmID -wwn SPportWWN add devs SymDevname |  
SymDevname -dir <#> -p <#>
```

For example:

```
symmask -sid 0128 -wwn 11:22:33:44:55:66:77:88 add devs 0014,0015  
-dir 16a -p 0
```

Preparing a Symmetrix system using ESN Manager

Note

To complete this procedure, you must be familiar with both Unisphere and ESN Manager.

Procedure

1. Use the native switch management tools to zone at least one port from an SP on the SAN Copy system to one or more FA ports on the participating Symmetrix systems (see Fibre Channel zoning requirements and recommendations in *EMC VNX for Block Concepts*).
-

Note

SAN Copy ports act as host initiators to remote systems (system without SAN Copy software).

2. Use ESN Manager version 2.1 to place the SAN Copy ports on the host side of the topology view.
 - If this is a new installation of the system, complete steps 2a and 2b.
 - If ESN Manager has already discovered the system and it is visible on the storage side of the topology view, complete all of step 2b.
-

Note

ESN Manager will not allow the SAN Copy ports (SP ports being used as initiators) to access the Symmetrix volumes unless the ports are visible on the host side of the topology view.

- a. Configure ESN Manager to discover the system.

Once ESN Manager discovers the system, it places an icon for it on the storage side of the topology view.
 - b. Use ESN Manager to move the system to the host side of the topology view.
 - a. In the Tools menu, select **Options > Out Band**.
 - b. In the **Out-Band** dialog box, select the system and clear the **Enable** checkbox.

Icons for the SAN Copy ports are displayed as hosts in the topology view.
3. Verify that the host with access to the Symmetrix volumes that are participating in the SAN Copy session is running the Unisphere Host Agent, and is either:
 - Connected to a system in the domain
or
 - Part of a portal configuration within the domain. For information about portals, see the online help or the *EMC Command Line Interface (CLI) Reference*.
 4. Register the SAN Copy ports with any other SP ports in the same zone using the `navicli -h sname sancopy -updateconnections` command (see [sancopy -updateconnections on page 521](#)).
- The SAN Copy ports log in to the Symmetrix storage ports.
5. To verify that the SAN Copy ports have been registered with the Symmetrix system, use the Symmetrix CLI command:

```
fpath lshist -d physicaldrive
```

where *physicaldrive* is the location of the VCM database on this host.

Note

Use the SymmCLI command `fpath lshostdev` to determine the pathname for the VCM database, or run `inq` on the host, if available.

6. Use ESN Manager to assign Symmetrix volumes to the SAN Copy ports.
-

Note

This is the same as assigning volumes to the host initiator.

7. Use ESN Manager to move the system back to the storage side of the topology view.
8. To verify that the SAN Copy port has access to the Symmetrix remote volumes, click the SAN Copy port entry and, in the volumes panel, make sure that ESN Manager surrounds any volumes that are accessible to the SP port with solid blue squares.

If you are copying data to a third-party system, use the third-party system management documentation to:

- Verify that the SAN Copy initiator has logged in to the system
- Provide access to the volume on the third-party system to the SAN Copy initiator

About admhost for SAN Copy

The `admhost` utility is an executable program that you can run interactively with a script. The `admhost` commands can activate and deactivate the source and destination logical units being used in a SAN Copy session, and can flush data from operating system buffers to ensure that the information on the source logical unit is current.

Note

Currently, `admhost` is supported only on the following Microsoft Windows platforms: Windows Server 2003, and Windows 2000. For the supported versions of these operating systems, see the *EMC SAN Copy and Admhost Release Notes*.

admhost SAN Copy command functions

This section lists the `admhost` commands and describes their function in SAN Copy sessions.

Note

The `admhost` utility is not supported on a UNIX host. Use the commands your operating system normally uses to discover new logical units, flush buffers, and unmount the logical units.

Table 47 admhost functions and descriptions

Function	Description	Syntax and examples
admhost lun_activate	<p>The admhost lun_activate command scans the bus for new devices and assigns a drive letter to all Symmetrix® and Block devices that do not already have one.</p> <p>Use the -l and -d switches to scan the bus for new devices and assign a drive letter to a device that has a specific World Wide Name (WWN). This command does not assign drive letters to any other devices.</p> <p>If admhost_activate is used and no new devices are found, and drive letters are already assigned to the Block and Symmetrix logical units, admhost reports an error message.</p>	<p>admhost lun_activate [-l lun_wwn -d drive_letter][-h] where</p> <p>-l lun_wwn specifies the World Wide Name (WWN) of the LUN to which you want to assign a drive letter (requires the -d switch).</p> <p>-d drive_letter specifies the drive letter you want to assign to the specified LUN (requires the -l switch).</p> <p>-h displays the help message.</p> <p>If the admhost lun_activate command activates more than one device, the drive letter is assigned to the first device in the list.</p>
admhost lun_deactivate	<p>The admhost lun_deactivate command attempts to take devices offline and make them inaccessible to the operating system. The command deletes the drive letter and dismounts the device.</p> <p>Use the admhost lun_activate command to bring the volume back online.</p> <p>Use this command on the source or destination host prior to starting a SAN Copy session (see Preparing logical units for a SAN Copy session on page 490).</p>	<p>admhost lun_deactivate -o object_name [-t object_type][-h] where</p> <p>-o object_name specifies the drive letter of the device you want to deactivate.</p> <p>When you specify the object name, admhost finds the device, deletes its drive letter, dismounts it, and then takes the device offline.</p> <p>-t object_type is an optional switch that specifies the object type (device or drive letter) of the device you want to deactivate.</p> <p>-h displays the help message.</p>
admhost lun_flush	<p>On a Windows host, the command flushes all data for the specified object_type so that the data in the write cache is written to disk.</p> <p>Use this command on the source or destination host prior to starting a SAN Copy session (see Preparing logical units for a SAN Copy session on page 490).</p>	<p>admhost lun_flush [-h] -o object_name [-t object_type device driveletter] where</p> <p>-h displays a help message.</p> <p>-o object_name specifies the name or drive letter of the device you want to flush.</p> <p>-t object_type is an optional switch that specifies the object type (device or drive letter) of the device you want to flush.</p>

Table 47 admhost functions and descriptions (continued)

Function	Description	Syntax and examples
		If you specify the device name, admhost identifies the corresponding device and drive letter.
admhost lun_list	Use to perform a list operation. This command lists the drive letters and WWNs for all host devices.	admhost lun_list [-h] [-l lun_wwn] [-d object_name] [-a output_style] where -h displays a help message. -l lun_wwn specifies the WWN of the device you want to list. -d object_name specifies the drive letter or name of the device you want to list. -a output_style specifies an alternative output style for the list output. The default style lists the drive letter and WWN for the logical unit. The -a switch lets you specify physical drive as the output style.
admhost help	This command displays the syntax for each admhost command and the revision number of the admhost utility. Output for the admhost help command is shown below. Revision number: V1.1.0.0.0 - 0 Valid operations include: lun_flush, lun_activate, lun_deactivate, lun_list, and help. Valid types include: device and driveletter.	admhost help

Preparing logical units for a SAN Copy session

Before you start a full SAN Copy session, and before you mark an incremental copy session, you must prepare the source and destination logical units that will be participating in the session. For a SAN Copy session with logical units connected to a Windows server, use the admhost utility. For all other types of servers, use the command specific to the server's operating system.

Use the following procedures to prepare source and destination logical units for full and incremental sessions.

Preparing source and destination logical units for full and incremental sessions

CAUTION

You must follow this procedure explicitly to avoid data loss or data corruption.

Procedure

1. Stop all I/O to the source and destination logical units.

Note

For incremental copy sessions, we strongly recommend that you stop all I/O to the source logical unit, and flush any server buffers prior to starting/marking the session. Once you receive confirmation that the mark or start session operation was successful, you can resume I/O to the source LUN.

If you do not stop all I/O and flush the server buffers to the source logical unit, the resultant data copy will be "crash consistent." This means that the contents of the copy will be similar to the contents left on a logical unit if the server crashes. You may be required to use the `fsck` or similar utility to make the data copy mountable. In this situation, the data copy may not contain the unflushed server data.

2. From the servers connected to any destination logical units, flush the host buffers and disconnect the logical units from the server by doing one of the following:

Note

If the destination logical units are not mounted, ignore this step.

- On a Windows Server 2003, or Windows 2000 host, use the following command to flush the host buffers and write the data to disk:
`admhost lun_flush -o destination drive letter`

For example, to flush the host buffers on drive F:, enter the following command:

```
admhost lun_flush -o F:
```

- On a Windows Server 2003 or Windows 2000 host, use the following command to take the disk offline:
`admhost lun_deactivate -o destination - drive letter`

For example, to take drive F: offline, enter the following command:

```
admhost lun_deactivate -o F:
```

- On a UNIX server, use the `umount` command to unmount the file system on any destination logical units.
 - On a NetWare server, use the `dismount` command to dismount the file system on any destination logical units.
3. From the server connected to the source logical units, deactivate source logical unit, flush the server buffers and disconnect the source logical units from the server by doing one of the following:

Note

If you want to maintain read access to the source logical unit, skip this step.

- On a Windows Server 2003 or Windows 2000 host, use the following command to flush the host buffers and write the data to disk:
`admhost lun_flush -o source -drive letter`

For example, to flush the host buffers on drive F:, enter the following command:

```
admhost lun_flush -o F:
```

- On a Windows Server 2003, or Windows 2000 host, use the following command to take the disk offline:
`admhost lun_deactivate -o source -drive letter`

For example, to take drive F: offline, enter the following command:

```
admhost lun_deactivate -o F:
```

- On a UNIX server, use the `umount` command to unmount the file system on any source logical units.
- On a NetWare server, use the `dismount` command to dismount the file system on any source logical units.

SAN Copy CLI commands

This section explains the `sancopy` commands used with `naviseccli` or `navicli` CLI commands to manage the SAN Copy software. These commands function only with a system that has the optional SAN Copy software installed.

SAN Copy command functions

The EMC VNX CLI for Block provides one of the two interfaces to the SAN Copy software. The Unisphere UI provides the other. The CLI command is `naviseccli sancopy` or `navicli sancopy` and you specify functions with switches.

The `sancopy` command functions, listed alphabetically, are:

- `sancopy -create`
- `sancopy -destinfo`
- `sancopy -duplicate`
- `sancopy -info`
- `sancopy -mark`
- `sancopy -modify`
- `sancopy -pause`
- `sancopy -remove`
- `sancopy -resume`
- `sancopy -settings`
- `sancopy -start`
- `sancopy -stop`
- `sancopy -throttle`
- `sancopy -transfer`
- `sancopy -unmark`
- `sancopy -updateconnections`

- sancopy -verify

The sancopy functions are listed functionally, in the order you would perform them, in the following table.

Table 48 navisecli or navicli sancopy command functions, listed in order of use

Command	Description
Essential functions (in order performed)	
sancopy -create, or sancopy -create -incremental	Creates a full SAN Copy session Creates an incremental SAN Copy session
sancopy -mark (incremental copy sessions only)	Creates a point-in-time copy of the changes that have taken place since the last copy session up to the time the session is marked
sancopy -start	Starts a SAN Copy session based on copy descriptor information
Optional status functions (alphabetically)	
sancopy -destinfo	Displays information on a particular session from the destination point of view
sancopy -info	Displays information on all SAN Copy sessions or displays attributes for a particular session
sancopy -stop	Stops a SAN Copy session
sancopy -verify	Ensures that the SAN Copy software can access a logical unit
Optional reconfiguration functions (alphabetically)	
sancopy -duplicate (full copy sessions only)	Creates a duplicate copy descriptor
sancopy -modify	Changes parameters in a copy descriptor
sancopy -pause	Pauses a SAN Copy session
sancopy -remove	Deletes a copy descriptor
sancopy -resume	Continues a paused, stopped or failed session
sancopy -throttle	Changes the throttle (execution rate) of an existing, running session

Table 48 naviseccli or navicli sancopy command functions, listed in order of use (continued)

Command	Description
sancopy -unmark (incremental copy sessions only)	Removes the existing mark in preparation for inserting a new mark
sancopy -transfer	Transfers control of a copy descriptor to the peer SP
sancopy -settings	Gets or sets environmental control parameters
sancopy -updateconnections	Updates system registration information

sancopy -create

Creates full and incremental SAN Copy sessions

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

Note

In addition to all the `sancopy -create` command switches, the `sancopy -create -incremental` command uses the following optional switches: `-changesonly`, `-linkbw`, and `-latency` (See Incremental SAN Copy sessions in *EMC VNX for Block Concepts*).

The `naviseccli` or `navicli sancopy` command with the `-create` switch creates a full SAN Copy session. A full session copies all the data on the source LUN to the destination logical units. When used with the optional `-incremental` switch, the `-create` command creates an incremental SAN Copy session. An incremental session copies only the data that has changed since the last copy session. This command supports thin LUNs as either sources or targets in the systems, if the system executing the CLI command (or the target system) is running a committed bundle that supports San Copy with thin LUNs.

Note

The destination logical unit(s) must exist when you execute this command.

By default a descriptor does not start a session unless you specify the `-start` switch.

For information on logical units accessible to the host use the basic CLI command `lunmapinfo` explained in the CLI reference manual.

SYNTAX

```
sancopy -create [-incremental] -name sessionName

-srcwwn source-lun-wwn | -srclun source -sp sourcelunnumber
[-device devicename]

-srcportwwn destportwwn destlunnumber

-destwwn dest-lun-wwn | -destlunnumber dest -sp destlunnumber|
-destportwwn destportwwn destlunnumber [-start]
[-throttle n]
[-verify]
[-o]
```

OPTIONS

-name *sessionName*

Specifies the SAN Copy session name. The name can be up to 64 characters and must be unique on a per-system basis. The names you create persist until you remove them.

For incremental SAN Copy sessions, SAN Copy creates a corresponding reserved session. The name assigned to the reserved session consists of the SAN Copy session name prefixed with the text "SANCopy_". This text is included in the 64-character limit. For example, if a SAN Copy session has the name *june_update*, the name of the corresponding reserved session is *SANCopy_june_update*.

-srcwwn *source-lun-wwn*

Note

For incremental sessions, the source LUN cannot be a SnapView snapshot, and it cannot be on a remote system. For incremental sessions, the source LUN must reside in the SAN Copy system.

Specifies the World Wide Name (WWN) of the source logical unit (LUN or snapshot) that will participate in the session. For full copy sessions, the source logical unit can reside in either the SAN Copy system or the remote system. If the source logical unit does not reside in the system specified by the *-h* switch, all destination logical units must reside in the system specified with *-h*.

By default, the connection type for the specified source logical unit will be set to *fibrepreferrred*. This sets the connection type to fibre or iSCSI (if fibre path could not be found) each time you start or resume a session. You can change the connection type using one of the following switches: *-fibre* or *-iscsi*.

Note

If the source logical unit does not reside in a remote system, setting the connection type will be ineffective.

-fibre

Sets the source logical unit's connection type to fibre.

-iscsi

Sets the source logical unit's connection type to iSCSI.

-srclun *source-sp* *sourcelunnumber*

An alternate way of adding a source logical unit entry to the copy descriptor. For *source-sp*, specify the SP hostname (or SP IP address) that owns the source logical unit. For *sourcelunnumber*, specify the logical unit number.

Note

For pre-FC4700 systems, if you use the `-srclun source-sp sourcelunnumber` function, you must use `-device devicename`.

By default, the connection type for the specified source logical unit will be set to `fibrepREFERRED`. This sets the connection type to fibre or iSCSI (if fibre path could not be found) each time you start or resume a session. You can change the connection type using one of the following switches: `-fibre` or `-iscsi`.

Note

If the source logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the source logical unit's connection type to fibre.

`-iscsi`

Sets the source logical unit's connection type to iSCSI.

`-device devicename`

Note

This command is required for pre-FC4700 systems.

Specifies the device name of the source LUN or destination LUN.

`-srcportwwn destportwwn destlunnumber`

Valid only for specifying source LUNs that reside in supported non-EMC systems. For *srcportwwn*, specify the World Wide Name of the SP port on the source system that owns the source LUN; for *srclunnumber* specify the number of the source LUN.

`-destwwn dest-lun-wwn`

Specifies which logical unit will be the destination of the copy. You can specify more than one destination logical unit if you want more than one copy of the source logical unit. If the destination logical units are not reachable by the client system, and you did not use the `-o` function, the CLI will display a warning and ask for confirmation.

By default, the connection type for the specified destination logical unit will be set to `fibrepREFERRED`. This sets the connection type to fibre or iSCSI (if a fibre path could not be found) each time you start or resume a session. You can change the connection type with one of the following switches: `-fibre` or `-iscsi`.

Note

If the destination logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the destination logical unit's connection type to fibre.

`-iscsi`

Sets the destination logical unit's connection type to iSCSI.

`-destlunnumber dest-sp destlunnumber`

An alternate way of adding a destination LUN entry to the copy descriptor. For *dest-sp*, specify the SP hostname (or SP IP address) that owns the destination logical unit; for *destlunnumber*, specify the LUN number. As with *-destwwn*, you can specify more than one destination logical unit if you want more than one copy of the source logical unit.

If you specify the *-verify* function, and the software cannot reach the SP or find the logical unit on the SP, it returns an error.

By default, the connection type for the specified destination logical unit will be set to *fibrepREFERRED*. This sets the connection type to fibre or iSCSI (if a fibre path could not be found) each time you start or resume a session. You can change the connection type with one of the following switches: *-fibre* or *-iscsi*.

Note

If the destination logical unit does not reside in a remote system, setting the connection type will be ineffective.

- fibre*
Sets the destination logical unit's connection type to fibre.
- iscsi*
Sets the destination logical unit's connection type to iSCSI.
- destportwwn destportwwn destlunnumber*
Valid only for specifying destination LUNs that reside in supported non-EMC systems. For *destportwwn*, specify the World Wide Name of the SP port on the destination system that owns the destination LUN; for *destlunnumber* specify the number of the destination LUN.
- start*
Starts a SAN Copy session using this copy descriptor. By default, the *-create* command does not start a session.
- throttle n*
Specifies the I/O rate of the copy. The maximum value is 10 and the minimum value is 1. If you omit this switch, the default value is 6.
- verify*
Verifies that SAN Copy can access all of the logical units specified in the *-create* command.
- o*
Suppresses a message asking you to confirm that you want to create a session now.
- incremental*
The following optional switches are used only with the *sancopy -create -incremental* command:
- changesonly*
If the source LUN and destination LUNs are synchronized when you create the session, use this switch. When you start the session, SAN Copy will copy only the changes that have occurred to the source LUN since the initial synchronization, and any subsequent sessions will copy only the changes.

If the source LUN and destination LUNs are not synchronized (or you are unsure of the copy state), do not use this switch. The first time you run the session, SAN Copy will copy the entire source LUN to the destination LUNs. Any subsequent sessions will copy only the changes.
- linkbw*
Specifies the available link bandwidth for the copy session. You can use the entire bandwidth of your network connection for incremental SAN Copy sessions, or you can specify only a portion of it.

`-latency`

The amount of time it takes to transmit a block of data to its destination expressed in milliseconds. Valid values are auto (SAN Copy computes the latency value), or a decimal number with no maximum limit. The default value is auto.

Note

If you create and start incremental sessions, and the copy is performed over a link that is slower than Fibre Channel, SAN Copy uses the link bandwidth and latency values to calculate the optimum buffer space for the session.

EXAMPLE

Full session example: For SP `ss1_spa`, the `-create` command creates a full copy session named `copy_lun66` that defines source logical unit 66 on this SP and destination logical unit 22 on host SP `ss5_spa`. The `-start` command then starts the copy session and the `-throttle` command specifies a throttle value of 3 for this session.

Incremental session example: The command below creates an incremental copy session named `copy_77` that defines logical unit 1 on SP `10.22.33.44` as the source and logical unit 22 on SP `ss5_spa` as the destination. The `-linkbw` command specifies a bandwidth value of 100 Mbits/sec., and the latency value defaults to auto.

```
naviseccli -h ss1_spa sancopy -create -name copy_lun66 -srclun
ss1_spa 66 -destlunnumber ss5_spa 22 - start -throttle 3
```

```
naviseccli -h 10.22.33.44 sancopy -create -incremental -name
copy_77 -srclun 10.22.33.44 1 -destlunnumber ss5_spa 22 -linkbw 100
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a command is not supported error message is printed to `stderr`. Other errors are:

```
DMF_NAME_SPECIFIED_NOT_UNIQUE
DMF_CREATE_PARAMS
DMF_SRCLUN_WWN_AND_NUMBER_SPECIFIED
DMF_DESTLUN_WWN_AND_NUMBER_SPECIFIED
DMF_INVALID_THROTTLE_VALUE
```

sancopy -destinfo

Displays information about destination logical units for a running session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

Note

In addition to all the `sancopy -destinfo` command switches listed above, incremental copy sessions use the optional `-syncstate` switch.

The `naviseccli` or `navicli sancopy` command with `-destinfo` gets session status from the destination point of view.

Note

If you have only full copy sessions, issuing the `-destinfo` command with no arguments displays all destination information for all copy sessions belonging to the specified SP.

If you have both full and incremental copy sessions, and you issue the `-destinfo` command with no arguments, some parameters for full sessions and all parameters specific to incremental sessions will not be displayed. In order to see all parameters for both full and incremental sessions, you must use the `-destinfo -all` command.

We recommend that you use the `-destinfo -all` command for both full and incremental sessions.

SYNTAX

```
sancopy-destinfo [-peer]
[-descriptorid descriptorid] | [-name name]
[-destwwn dest-lun-wwn] [-destlunnumber destsp destlunnumber]
[-destportwwn destportwwndestlunnumber]
[-complete]
[-counts]
[-failure]
[-time]
[-connectiontype]
[-all]
```

OPTIONS

`-peer`

Lists all destination information for the peer SP.

`-descriptorid descriptorid | -name name`

Used alone, `-descriptorid` or `-name` displays descriptor IDs or names of all copy descriptors of the specified SP; `-descriptorid` with an ID or `-name` with a name lists the status (WWN, failure status, transfer count, percent complete and failure time) of all destinations of the copy descriptor.

`-destwwn dest-lun-wwn`

Used alone, `-destwwn` displays the WWNs of all destinations of the descriptors. `-destwwn` with a WWN lists the status of the specified destination logical unit.

`-destlunnumber destsp destlunnumber`

The switch `-destlunnumber` cannot be used alone. When used with the SP name and LUN number, it displays the status of the specified destination logical unit.

`-destportwwn destportwwn lunnumber`

The switch `-destportwwn` displays the status of the specified destination LUN that resides in a supported non-EMC system.

`-complete`

Displays the percentage complete of the copy session when any destination logical unit(s) failed. Displays N/A when there is no failure.

`-counts`

Displays the number of disk blocks transferred when any destination logical units failed. Displays N/A when there is no failure.

`-failure`

Displays failure status or `No Failure` for destination logical unit(s). For more information, also use the `-time`, `-counts`, and/or `-complete` switches.

`-time`

Displays the time when any destination logical unit(s) failed. Displays N/A when there is no failure.

`-connectiontype`

Displays the requested and the actual connection type used for all destination logical units participating in the copy session.

Valid values for Requested Connection Type are fibre, iSCSI or fibre preferred. Valid values for Actual Connection Type are fibre, iSCSI, undetermined and N/A (for destination logical units within the same system).

`-all`

Displays all destination information for both full and incremental copy sessions belonging to the specified SP.

Note

If you have both full and incremental copy sessions, and you issue the `-destinfo` command with no arguments, some parameters for full sessions and all parameters specific to incremental sessions will not be displayed. In order to see all parameters for both full and incremental sessions, use the `destinfo -all` command.

The following optional switch is used only with incremental sessions:

`-syncstate`

Displays the state of the data of the specified destination logical unit with respect to the source logical unit. Valid states are N/A, Inconsistent, and Initial Copy Required.

EXAMPLE

This example displays the output for a full copy session using the `-destinfo` command without the `-all` switch. In order to display all the information for an incremental copy session, use the `-destinfo -all` command.

```
naviseccli -h dm340c-16a sancopy -destinfo -connection type
```

```
Copy Descriptor Name: Std0_1
Copy Descriptor ID: 7202
Owner: SPA

Requested Connection Type: Fibre Preferred
Actual Connection Type: N/A

Requested Connection Type: Fibre Preferred
Actual Connection Type: N/A
```

In the example below, the session includes two destination logical units. The requested connection type and the actual connection type are displayed for both destinations.

```
naviseccli -h ssl_spa sancopy -destinfo
```

```
Copy Descriptor Name: DMHOST7_SAN005_10.6.92.169_00_p000
Copy Descriptor ID: 8
Owner: SPB
Copy Descriptor Type: Incremental
Copy Whole LUN: NO
```

```

Source LUN WWN: 60:06:01:60:0E:C2:1C:00:84:C4:31:B2:1C:24:DC:11
Source PortWWN and LUN: N/A
Number of Blocks to Copy: 0
Size of source LUN: 8388608
Number Of Destinations: 2
Destination LUN WWN:
    60:06:01:60:61:31:19:00:FD:C8:6D:54:1C:24:DC:11,
    60:06:01:60:61:31:19:00:6E:F0:C3:61:1C:24:DC:11
Destination PortWWN and LUN: N/A
Session Status: Complete
Initial Throttle: 4
Current Throttle: 4
Transfer Count: 8388608
Percent Complete: 100
Start Time: 06/27/07 10:51:05
Completion Time: 06/27/07 12:55:12
Duration: 2 hr 4 min 7 sec
Failure Status: No Failure
Marked: NO
Mark Time: N/A
Blocks modified since mark: N/A
Link Bandwidth (Mbps): 1.024
Effective Bandwidth (Mbps): 0.719
Latency (msec): Auto
Measured Latency (msec): 6.140
Requested Connection Type: Fibre Preferred
Actual Connection Type: N/A

```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Command is not supported error message is printed to stderr.

sancopy -duplicate

Makes a copy of a copy descriptor

PREREQUISITES

For navisecli, you must have a user account on the system on which you want to execute the command.

For navicli, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The navisecli or navicli sancopy command with `-duplicate` makes a copy of an existing copy descriptor. You can then modify the copy as needed (`-modify`). The copy descriptor that you want to duplicate can be part of an active SAN Copy session. You do not need to stop the copy descriptor before duplicating it.

Note

The sancopy `-duplicate` command is not supported for incremental SAN Copy sessions.

SYNTAX

```
sancopy -duplicate -descriptorid descriptorid|-name name
```

OPTIONS

`-descriptorid` *descriptorid*|-name *name*
Identifies the descriptor you want to duplicate.

EXAMPLE

This command duplicates SP `ss1_spa`'s copy descriptor `copy_lun_66` and assigns a default name to the new descriptor - copy of `xxx`, where `xxx` is the name of the descriptor you are duplicating, for example, a copy of `copy_lun_66`.

```
naviseccli -h ss1_spa sancopy -duplicate -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

sancopy -info

Displays information about copy descriptors

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a `username/host` entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION**Note**

In addition to all the `sancopy -info` command switches listed above, incremental copy sessions use the following optional switches: `-copywholelun`, `-markstatus`, `-srcsize`, `-link` (see Incremental SAN Copy sessions in *EMC VNX for Block Concepts*).

The `naviseccli` or `navicli sancopy` command with `-info` displays information on copy sessions. It differs from `-destinfo` in that `-destinfo` gets information on the destination logical units.

Note

If you have only full copy sessions, issuing the `-info` command with no arguments displays all destination information for all copy sessions belonging to the specified SP.

If you have both full and incremental copy sessions, and you issue the `-info` command with no arguments, some parameters for full sessions and all parameters specific to incremental sessions will not be displayed. In order to see all parameters for both full and incremental sessions, use the `-info -all` command.

We recommend that you use the `-info -all` command for both full and incremental sessions.

SYNTAX

```
sancopy -info [-peer]
[-descriptorid descriptorid|-name name]
```

```

[-srcwwn]
[-length]
[-numdests]
[-dests]
[-sessionstatus]
[-throttle]
[-counts]
[-complete]
[-time]
[-failure]
[-duration]
[-type]
[-connectiontype]
[-all]

```

OPTIONS

`-peer`

Lists all destination information on the peer SP.

`-descriptorid descriptorid`

Lists the descriptor IDs of all sessions belonging to the specified SP when you issue this command without a specific descriptor ID.

Lists the status of the specified session when you issue this command with a specific descriptor ID.

`-name name`

Lists the names of all copy descriptors belonging to the specified SP when you issue this command without a specific copy descriptor name.

Lists the status of the specified copy descriptor when you issue this command with a name.

`-srcwwn`

Displays the WWNs of all source logical units.

`-length`

Displays the number of blocks to copy. Displays N/A if the descriptor was created using the source WWN.

`-numdests`

Displays the number of destination logical units.

`-dests`

Displays the WWNs of all destinations logical units in the specified descriptor.

`-sessionstatus`

Displays the status of the session associated with the copy descriptor: not started, complete, stop, pause, or failed.

`-throttle`

Displays the current throttle value for the copy session.

`-counts`

Displays the number of disk blocks transferred.

`-complete`

Displays the percentage complete of the copy session.

`-time`

Displays the time when the copy completes.

`-failure`

Displays failure status or `No Failure` for the copy descriptor's source and destination logical units.

`-duration`

Displays the completion time and duration of the session.

`-type`

- Displays the session type - Full or Incremental.
- `-connectiontype`
Displays the requested and the actual connection type used for the source logical unit participating in the copy session.

Valid values for Requested Connection Type are fibre, iSCSI or fibre preferred. Valid values for Actual Connection Type are fibre, iSCSI, undetermined and N/A (for source logical units within the same system).
- `-copywholelun`
Displays whether the session will copy the entire source logical unit or only changes to it. If Yes, the session copies the entire logical unit; if No, it copies only the changes.
- `-markstatus`

Note

If a session is marked, SAN Copy creates a point-in-time copy of the changes that have taken place since the last copy session up to the time the session is marked.

Displays whether the session has been marked, and if yes, displays the time of the mark, and the number of blocks of data that have changed since the mark occurred (these data blocks will not be copied when the session starts).

- `-srcsize`
Displays the size, in blocks, of the source logical unit.
- `-link`
Displays the status of any network link utilization performance parameters as follows:

Link Bandwidth — Link bandwidth of the network connection set when the session was created. Value is in Mb/sec.

Effective Bandwidth — Actual link bandwidth for the session measured by SAN Copy.

Latency — Displays the latency value set when the session was created or the default value of Auto. Latency is the amount of time it takes to transmit a block of data to its destination expressed in milliseconds.

Measured Latency — The actual measured latency of the session expressed in milliseconds.
- `-all`
Displays all session information for both full and incremental copy sessions belonging to the specified SP.

Note

You must use the `-destinfo -all` command to display information that is specific to incremental sessions. We recommend that you use the `-destinfo -all` command for both full and incremental sessions.

The information displayed with the `-all` command may change in future revisions of the software.

EXAMPLE

This command displays the status of all copy descriptors for both SPs, but does not include any information specific to incremental copy sessions. The format is as follows:

naviseccli sancopy -info -all

```

Copy Descriptor name: session1
Copy Descriptor ID: 4
Owner: SP A
Copy Descriptor Type: Full
Source LUN WWN: 60:06:01:60:72:70:18:00:5A:FE:9A:58:0E:5A:DB:11
Source PortWWN and LUN: N/A
Number of Blocks to Copy: 4194304
Number Of Destinations: 1
Destination LUN WWN: 60:06:01:60:72:70:18:00:5B:FE:9A:58:0E:5A:DB:11
Destination PortWWN and LUN: N/A
Session Status: A session was active when the Storage Processor (SP)
rebooted. The session may be restarted.
Initial Throttle: 6
Current Throttle: 4
Transfer Count: 158720
Percent Complete: 3
Start Time: 10/12/06 16:54:57
Completion Time: 10/12/06 16:59:44
Duration: 4 min. 47 sec
Failure Status: N/A

Copy Descriptor name: ses_change_ar
Copy Descriptor ID: 1
Owner: SP B
Copy Descriptor Type: Full
Source LUN WWN: 60:06:01:60:72:70:18:00:C4:AD:19:66:0E:5A:DB:11
Source PortWWN and LUN: N/A
Number of Blocks to Copy: 4194304
Number Of Destinations: 1
Destination LUN WWN: 60:06:01:60:72:70:18:00:C45AD:19:66:0E:5A:DB:11
Destination PortWWN and LUN: N/A
Session Status: Session is paused.
Initial Throttle: 6
Current Throttle: 6
Transfer Count: 799744
Percent Complete: 19
Start Time: 10/12/06 17:17:37
Completion Time: N/A
Duration: N/A
Failure Status: No Failure

Copy Descriptor name: SAN000_10.6.84.185_00_p1608
Copy Descriptor ID: 21646
Copy Descriptor Type: Full
Source LUN WWN: 60:06:01:60:E3:FA:15:00:0E:4E:70:A1:8B:C5:DA:11
Source PortWWN and LUN: 60:06:01:60:E3:FA:15:00:0E:4E:
70:A1:8B:C5:DA:11
Number of Blocks to Copy: 2097152
Number Of Destinations: 1
Destination LUN WWN: 60:06:01:60:E3:FA:15:00:0F:4E:70:A1:8B:C5:DA:11
Destination PortWWN and LUN: N/A
Session Status: Complete
Initial Throttle: 4
Current Throttle: 4
Transfer Count: 2097152
Percent Complete: 100
Start Time: 04/19/06 09:55:39
Completion Time: 04/19/06 09:55:57
Duration: 18 sec
Failure Status: No Failure

```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a Command is not supported error message is printed to stderr.

sancopy -mark

For incremental sessions only--Makes a point-in-time copy of changes to the source logical unit since the last copy session.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-mark` creates a point-in-time copy of the changes that have taken place since the last copy session up to the time the session is marked. You can mark the session at any time before you start the session, or SAN Copy will mark the session at the time the session starts. When you start a marked session, all changes made to the source LUN up to the time of the mark will be copied to the destination logical units.

Note

Before marking the session, make sure that you prepare the source logical unit as described in the section, [Preparing logical units for a SAN Copy session on page 490](#). Once you receive confirmation that the mark operation was successful, you can resume I/O to the source logical unit.

SYNTAX

```
sancopy -mark -descriptorid descriptorid | -name sessionName
```

OPTIONS

```
-descriptorid descriptorid | -name sessionName
```

Identifies the session you want to mark.

EXAMPLE

This command marks the SAN Copy session named `monday`. When the session is started, it copies only the changes to the source logical unit that have taken place since the last copy session up to the time of the mark.

```
naviseccli -h 10.11.22.33 sancopy -mark -name monday
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

sancopy -modify

Changes the definitions for a copy session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli sancopy` command with `-modify` changes an existing SAN Copy session. This command will fail if the session has started or is in the paused state. You must stop the session before modifying the session. This command supports thin LUNs only when the system executing the CLI command (or the target system) is running a committed bundle that supports San Copy with thin LUNs.

The new session information will not take effect until you start a new session.

You must include a copy descriptor ID or name and one or more switches to indicate changes.

Note

If you are modifying an incremental SAN Copy session, the following optional switches are available: `-copywholelun`, `-chglinkbw`, and `-chglatency`.

SYNTAX

```
sancopy -modify -descriptorid descriptor-id|-name name
[-addestbywwn dest-lun-wwn]
[-addestbylunnumber destsp destlunnumber]
[-addestbyportwwn destportwwn destlunnumber]
[-chgdestconnectiontypebywwn destwwn]
[-chgdestconnectiontypebylunnumber destspname destlunnumber]
[-chgname newname]
[-chgsrbylunnumber srcspname srclunnumber]
[-chgsrbywwn srcwwn]
[-chgsrbyportwwn srcportwwn srclunnumber]
[-chgsrconnectiontype srcwwn]
[-chgsessiontype]
[-o]
[-rmdestbylunnumber destsp destlunnumber]
[-rmdestbywwn dest-lun-wwn]
[-rmdestbyportwwn destportwwn destlunnumber]
[-throttle n]
```

OPTIONS

`-descriptorid` *descriptor-id* | `-name` *name*

Identifies the descriptor you want to change.

`-addestbywwn` *dest-lun-wwn*

Adds a destination logical unit to the descriptor by World Wide Name.

By default, the connection type for the specified logical unit will be set to `fibrepreferred`. This sets the connection type to fibre or iSCSI (if fibre path

could not be found) each time you start or resume a session. You can change the connection type using one of the following switches: `-fibre` or `-iscsi`.

Note

If the logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the logical unit's connection type to fibre.

`-iscsi`

Sets the logical unit's connection type to iSCSI.

`-adddestbylunnumber destsp destlunnumber`

Adds a destination logical unit to the descriptor by destination SP (specify the SP name or IP address) and destination logical unit number.

By default, the connection type for the specified logical unit will be set to `fibrepreferred`. This sets the connection type to fibre or iSCSI (if fibre path could not be found) each time you start or resume a session. You can change the connection type with one of the following switches: `-fibre` or `-iscsi`.

Note

If the logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the logical unit's connection type to fibre.

`-iscsi`

Sets the logical unit's connection type to iSCSI.

`-adddestbyportwwn destportwwn destlunnumber`

Adds only destination logical units that reside in supported non-EMC systems. For *destportwwn*, specify the World Wide Name of the SP port on the destination system that owns the destination LUN; for *destlunnumber* specify the number of the destination LUN.

`-chgdestconnectiontypebywwn destwwn`

Changes the connection type of the destination logical unit using its World Wide Name, *destwwn*.

You can change the connection type of the logical unit using one of the following switches: `-fibre`, `-iscsi`, or `-fibrepreferred`.

Note

If the destination logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the destination logical unit's connection type to fibre.

`-iscsi`

Sets the destination logical unit's connection type to iSCSI.

`-fibrepreferred`

Sets the destination logical unit's connection type to fibre or iSCSI (if fibre path could not be found).

`-chgdestconnectiontypebylunnumber destspname destlunnumber`

Changes the connection type of the destination logical unit using its destination SP (specify the SP name or IP address), *destspname*, and its destination logical unit number, *destlunnumber*.

You can change the connection type of the logical unit with one of the following switches: `-fibre`, `-iscsi`, or `-fibrepreferred`.

Note

If the destination logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the destination logical unit's connection type to fibre.

`-iscsi`

Sets the destination logical unit's connection type to iSCSI.

`-fibrepreferred`

Sets the destination logical unit's connection type to fibre or iSCSI (if a fibre path could not be found).

`-chgname newname`

Note

Not available for incremental sessions.

Changes the descriptor name to *newname*.

`-chgsrcbylunnumber srcspname srclunnumber`

Note

Not available for incremental sessions.

Changes the source logical unit in the session owned by source SP (for *srcspname*, specify the SP name or IP address) and source logical unit number.

By default, the connection type for the specified source logical unit will be set to `fibrepreferred`. This sets the connection type to fibre or iSCSI (if fibre path could not be found) each time you start or resume a session. You can change the connection type with one of the following switches: `-fibre` or `-iscsi`.

Note

If the source logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the source logical unit's connection type to fibre.

`-iscsi`

Sets the source logical unit's connection type to iSCSI.

`-chgsrcbywwn srcwwn`

Note

Not available for incremental sessions.

Changes the source logical unit in the descriptor using the World Wide Name of the new source logical unit.

By default, the connection type for the specified source logical unit will be set to `fibrepreferred`. This sets the connection type to fibre or iSCSI (if a fibre path could not be found) each time you start or resume a session. You can change the connection type with one of the following switches: `-fibre` or `-iscsi`.

Note

If the source logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the source logical unit's connection type to fibre.

`-iscsi`

Sets the source logical unit's connection type to iSCSI.

`-chgsrbyportwwn srcportwwn srclunnumber`

Valid only for changing source logical units that reside in supported non-EMC systems. For *srcportwwn*, specify the World Wide Name of the SP port on the source system that owns the source LUN; for *srclunnumber* specify the number of the source LUN.

`-chgsrconnectiontype srcwwn`

Changes the connection type of the source logical unit using its World Wide Name (*srcwwn*).

You can change the connection type of the logical unit with one of the following switches: `-fibre`, `-iscsi`, or `-fibrepreferred`.

Note

If the logical unit does not reside in a remote system, setting the connection type will be ineffective.

`-fibre`

Sets the session connection type to fibre.

`-iscsi`

Sets the session connection type to iSCSI.

`-fibrepreferred`

For connections between combo systems (Fibre Channel and iSCSI), sets the connection type to fibre, if possible. If not possible, uses iSCSI.

`-chgsessiontype`

Changes the type of the SAN Copy session — incremental or full.

If you change the session type from full to incremental, the following optional switches are available: `-changesonly`, `-linkbw`, and `-latency`.

`-changesonly`

If the source LUN and destination LUNs are synchronized, use this switch. When you start the session, SAN Copy will copy only the changes that have occurred to the source LUN since the initial synchronization, and any subsequent sessions will copy only the changes.

If the source LUN and destination LUNs are not synchronized (or you are unsure of the copy state), do not use this switch. The first time you run the session, SAN Copy will copy the entire source LUN to the destination LUNs. Any subsequent sessions will copy only the changes.

`-linkbw`

Specifies the available link bandwidth for the copy session. You can use the entire bandwidth of your network connection for incremental SAN Copy sessions, or you can specify only a portion of it. The default value is .15 mbps (equivalent to a T1 line).

`-latency`

Specifies the latency value for the incremental session. Latency is the amount of time it takes to transmit a block of data to its destination expressed in milliseconds. Valid values are auto (SAN Copy computes the latency value), or a decimal number

with no maximum limit. The default value is auto. Entering a 0 latency value is the equivalent to entering auto.

Note

If you create and start incremental sessions, and the copy is performed over a link that is slower than Fibre Channel, SAN Copy uses the link bandwidth and latency values to calculate the optimum buffer space for the session.

- o Suppresses a confirmation message describing the descriptor.
- rmdesbylunnumber *destsp destlunnumber*
Removes a destination logical unit from the specified descriptor by destination SP (specify the SP name or IP address) and destination logical unit number.
- rmdesbywwn *dest-lun-wwn*
Removes a destination logical unit from the specified descriptor by World Wide Name.
- rmdesbyportwwn *destportwwn destlunnumber*
Valid only for removing destination logical units that reside in supported non-EMC systems. For *destportwwn*, specify the World Wide Name of the SP port on the destination system that owns the destination LUN; for *destlunnumber* specify the number of the destination LUN.
- throttle n
Changes the I/O rate to n (valid values 1 through 10). The default value is 6. To change the throttle value of an active session, use the command `sancopy -throttle`.
For incremental SAN Copy sessions, the following optional switches are available.
- copywholelun
Specifies that SAN Copy copy the entire contents of the source logical unit the next time the session is run.
- chglinkbw
Changes the link bandwidth for the specified session.
The default value is .15 mbps (equivalent to a T1 line).
- chglatency
Changes the latency value for the specified session. Latency is the amount of time it takes to transmit a block of data to its destination expressed in milliseconds. Valid values are auto (SAN Copy computes the latency value), or a decimal number with no maximum limit. The default value is auto.

EXAMPLE

For SP `ss1_spa`, copy descriptor `copy_66` is a full copy session. This command adds the destination logical unit whose World Wide Name is `60:06:01:6B:61:09:00:00:6A:6F:9F:ED:0B:A3:D6:11`. For SP `ss1_spa`, copy descriptor `copy_77` is an incremental copy session. The next time session `copy_77` is run, SAN Copy will copy the contents of the entire source LUN to the destinations.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa sancopy -modify -name copy_66 -adddestbywwn
60:06:01:6B:61:09:00:00:6A:6F:9F:ED:0B:A3:D6:11
```

```
naviseccli -h ss1_spa sancopy -modify -name copy_77 -copywholelun
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
DMF_INVALID_THROTTLE_VALUE
DMF_REMOVED_DEST_DOESNOT_EXIST
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
SESSION_NAME_EXISTS
```

sancopy -pause

Suspends I/O in an active SAN Copy session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a `username/host` entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-pause` pauses an existing SAN Copy session. You must specify the descriptor ID or name.

To continue the paused session, use the `sancopy -resume` function.

SYNTAX

```
sancopy -pause -descriptorid descriptor -id|-name name
```

OPTIONS

```
-descriptorid descriptor-id|-name name
```

Identifies the descriptor whose session you want to pause.

EXAMPLE

For SP `ssl_spa`, this command pauses the session with the `copy_lun_66` descriptor.

```
naviseccli -h ssl_spa sancopy -pause -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```


Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
INVALID_SIZE
```

sancopy -remove

Deletes a copy session.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-remove` deletes one or all previously stored copy sessions. A copy session cannot be removed if it is currently active. An active copy session must complete the copy first or be stopped.

If you omit the override `-o` option, the CLI displays a message asking you to confirm the removal of the copy session.

This command can use only one of the following switches at a time:

```
-descriptorid; -name; -all
```

SYNTAX

```
sancopy -remove -descriptorid descriptor-id|-name name|-all [-o]
```

OPTIONS

```
-descriptorid descriptor-id|-name name
```

Identifies the descriptor you want to remove.

```
-all
```

Directs the CLI to remove all descriptors on the SP.

```
-o
```

Suppresses a message confirming the removal of the descriptor.

EXAMPLE

For SP `ss1_spa`, this command deletes the copy descriptor named `copy_lun_80`.

```
naviseccli -h ss1_spa sancopy -remove -name copy_lun_80
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
Do you really want to remove the copy descriptor(y/n) [n]? y
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
```

sancopy -resume

Continues I/O in a SAN Copy session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-resume` resumes an existing SAN Copy session that was paused, stopped, or failed. If the copy session was paused, it will start at the paused point.

If the copy session was stopped or failed, the `-resume` command will restart the copy session from the latest checkpoint (explained in the `settings` command). If the checkpoint feature is disabled (the checkpoint was set to `disable`), the copy session will restart from the beginning.

You must specify the descriptor ID or name.

SYNTAX

```
sancopy -resume -descriptorid descriptor-id | -name name
```

OPTIONS

```
-descriptorid descriptor-id | -name name
```

Identifies the descriptor whose session you want to resume.

EXAMPLE

For SP `ss1_spa`, this command resumes the `sancopy_lun_66` copy session.

```
naviseccli -h ss1_spa sancopy -resume -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other error messages are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
SESSION_ALREADY_IN_PROGRESS
```

sancopy -settings

Lists or changes SAN Copy environmental parameters

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-settings` gets and sets four SAN Copy environmental control parameters :

- the maximum number of concurrent active sessions
- a checkpoint interval for session status
- number of buffers per session
- buffer size

To find the current settings, use the `-list` switch; to change a setting, omit `-list` and specify the setting.

If you do not change the setting specifications, the software uses default values.

SYNTAX

```
sancopy -settings -list
```

```
sancopy -settings [-setmax MaxNumSessions]
[-setcheckpoint CheckpointInterval]
[-setbufsize buffersize]
[-setbufnum buffervalue]
```

OPTIONS

`-list`

Displays the current environmental parameter settings.

`-setmax maxNumSessions`

Sets the maximum number of SAN Copy sessions. The maximum number of sessions depends on the system type.

`-setcheckpoint CheckpointInterval`

Sets the checkpoint interval in minutes (values 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, 55, 60, 65, 70, 75, 80, 85, or 90) or, with the literal `Disable`, disables the checkpoint feature. By default, checkpointing is enabled with a value of 10 minutes.

The SAN Copy software can track the progress of a copy session by checkpoint interval. If a session fails while checkpointing is enabled, then after correcting the problem, you can direct the software to resume (`-resume` function). It will continue copying from the last checkpoint, if possible. If checkpointing is disabled, or if resuming is not possible, you can start copying from the beginning.

`-setbufsize` *buffersize*

Specifies the maximum buffer size to be used by SAN Copy sessions on a Fibre Channel connection. The default buffer size is 1024 blocks. The minimum buffer size is 64 KB.

`-setbufnum` *buffernumber*

Specifies the maximum number of buffers to be used by SAN Copy sessions on a Fibre Channel connection. The default number is 4, and the minimum is 1.

EXAMPLE

For SP `ssl_spa`, this command lists the current value for maximum number of sessions, checkpoint interval, number of buffers per session, and buffer size. For SP `ssl_spa`, this command changes the checkpoint interval to 45 minutes. For SP `ssl_spa`, this command changes the number of buffers to 3.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ssl_spa sancopy -settings -list
```

```
naviseccli -h ssl_spa sancopy -settings -setcheckpoint 45
```

```
naviseccli -h ssl_spa sancopy -settings -setbufnum 3
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

Other errors are:

```
DMF_LIST_AND_SET_SPECIFIED
DMF_INVALID_MAX_SESSION_VALUE
DMF_INVALID_CHECKPOINT_INTERVAL_VALUE
```

Errors returned from driver or agent are:

```
SESSION_MAX_OVER_LIMIT
SESSION_POLL_OVER_LIMIT
```

sancopy -start

Starts a SAN Copy session based on values in a copy descriptor

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-start` starts a SAN Copy session with the source and destination logical unit(s) identified in the copy descriptor. You can start a session using the descriptor ID or name.

The command works only when issued to the SP that owns the copy descriptor.

The `sancopy -info` command provides useful information for the `sancopy -start` command.

Note

Before starting a copy session, make sure you prepare the source and destination logical units as described in the section, [Preparing logical units for a SAN Copy session on page 490](#). When you start a SAN Copy session, SAN Copy automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
sancopy -start -descriptorid descriptor-id | -name name
[-copywholelun [-nomark new | all] ]
```

OPTIONS

`-descriptorid` *descriptor-id* | `-name` *name*

Identifies the copy descriptor for the session you want to start.

For incremental SAN Copy sessions, the following optional switches are available.

`-copywholelun`

Specifies that this incremental session will copy the entire contents of the source logical unit to all the destination logical units. To use this command the session must be unmarked. When you start the session using this command, SAN Copy will mark the session for you. This command overrides any existing settings for the session.

`-nomark` *new|all*

Used with the `-copywholelun` switch, and works only if the session is not already marked. Specifies that this session will not be marked when it is started. A session run with the `-copywholelun -nomark` command completes more quickly, but at the end of the session, the source and destination logical units are inconsistent. To make them consistent, you must run the session again.

When `-nomark` is used with the *new* option, SAN Copy starts an unmarked session that copies the contents of the source logical unit to any new destinations.

When `-nomark` is used with the *all* option, SAN Copy starts an unmarked session that copies the contents of the source logical unit to all destinations.

EXAMPLE

For SP `ss1_spa`, this command starts a copy session based on the copy descriptor named `copy_lun_66`.

```
naviseccli -h ss1_spa sancopy -start -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the VNX OE for Block software running in the SP does not support the command, the CLI prints a `Command is not supported` error message to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
SESSION_ALREADY_IN_PROGRESS
```

sancopy -stop

Stops (terminates) a SAN Copy session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command `-stop` terminates the specified session. The session stops automatically when it is complete, so this command is not needed in the normal course of SAN Copy operations.

SYNTAX

```
sancopy -stop -descriptorid descriptor-id | -name name
```

OPTIONS

```
-descriptorid descriptor-id | -name name
```

Identifies the descriptor whose session you want to stop.

EXAMPLE

For SP `ssl_spa`, this command stops the copy session with the copy descriptor name `copy_lun_66`.

```
naviseccli -h ssl_spa sancopy -stop -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other error messages are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
```

sancopy -throttle

Changes the I/O rate of an active copy descriptor

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-throttle` lets you change the throttle of an existing, running copy descriptor. The new throttle value is valid only while the session is running and will not change the value in the descriptor database.

To change the throttle value of an existing copy descriptor that is not running, use the `sancopy -modify -throttle` command (see [sancopy -modify on page 507](#))

SYNTAX

```
sancopy -throttle -descriptorid descriptor-id | -name name n
```

OPTIONS

`-descriptorid descriptor-id` | `-name name`

Identifies the descriptor whose throttle value you want to change.

`n`

Is the new throttle value, ranging from 1 (slowest) to 10 (highest). The default value is 6.

EXAMPLE

For SP `ss1_spa` descriptor `copy_lun_66`, this command displays information, including throttle information. For SP `ss1_spa` descriptor `copy_lun_66`, these commands get information (including throttle value), then change the throttle value to 8.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa sancopy -info -name copy_lun_66
```

```
naviseccli -h ss1_spa sancopy -throttle -name copy_lun_66 8
```

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_INVALID_THROTTLE_VALUE
DMF_UNKNOWN_DESCRIPTOR
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
```

sancopy -transfer

Transfers a copy descriptor from the peer SP

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-transfer` moves the specified descriptor from the peer SP to the SP you specify in the command.

After transferring the copy descriptor, you can start a new session on the working SP using that descriptor.

SYNTAX

```
sancopy -transfer -descriptorid descriptor-id | -name name
```

OPTIONS

```
-descriptorid descriptor-id | -name name
```

Identifies the descriptor you want to transfer to the working SP.

EXAMPLE

This command transfers the copy descriptor `copy_lun_66` to the SP named `ss1_spb`.

```
naviseccli -h ss1_spb sancopy -transfer -name copy_lun_66
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED
DMF_NAME_OR_UID_REQUIRED
DMF_UNKNOWN_DESCRIPTOR
DMF_PEER_SP_ALIVE
```

Errors returned from the driver or agent are:

```
UNKNOWN_DESCRIPTOR_ID
UNKNOWN_DESCRIPTOR_NAME
```


sancopy -unmark

For incremental sessions only--Removes the point-in-time copy from a marked incremental SAN Copy session

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-unmark` removes the point-in-time copy from a marked incremental SAN Copy session.

Note

You cannot unmark a session if any of the session's destination logical units have failed.

SYNTAX

```
sancopy -unmark -descriptorid descriptor-id | -name name
```

OPTIONS

```
-descriptorid descriptor-id | -name name  
    Identifies the session you want to unmark.
```

EXAMPLE

This command removes the mark from the incremental copy session named `monday`.

```
naviseccli -h 10.11.22.33 sancopy -unmark -name monday
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

If the version of VNX OE for Block software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors are:

```
DMF_NAME_AND_UID_SPECIFIED  
DMF_NAME_OR_UID_REQUIRED  
DMF_UNKNOWN_DESCRIPTOR
```

sancopy -updateconnections

Updates connection and registration information

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-updateconnections` updates registration information for all attached systems.

You should run this command for each SP.

SYNTAX

```
sancopy -updateconnections
```

EXAMPLE

These commands update connections on both SPs, `ssl_spa` and `ssl_spb`.

```
naviseccli -h ssl_spa sancopy -updateconnections
```

```
naviseccli -h ssl_spb sancopy -updateconnections
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

sancopy -verify

Ensures that the SAN Copy software can access a logical unit

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command.

For `navicli`, you must have a username/host entry in the Unisphere SP agent configuration file to use this command.

DESCRIPTION

The `naviseccli` or `navicli` `sancopy` command with `-verify` ensures that the SAN Copy software can access a logical unit. You should issue the command to both SPs as one SP might have access, where the other does not.

SYNTAX

```
sancopy -verify -wwn lun-wwn | -lunnumber sp-name lunnumber
[-device devicename]
[-log]
[-iscsi|-fibre]
```

OPTIONS

`-wwn lun-wwn`

Specifies the World Wide Name of the logical units. You can specify more than one WWN.

`-lunnumber sp-name lunnumber`

An alternate way of specifying a logical unit. For `sp-name`, specify the SP hostname (or SP IP address). For `lunnumber`, specify the source or destination logical unit number. You can specify more than one LUN number.

`-device devicename`

Specifies the system device name.

`-log`

Logs events to the SP event log.

Note

If SAN Copy cannot find a LUN on a remote system, use these events to help determine the reason.

`-iscsi|-fibre`

Specifies which connection type to use to verify the connection. If a connection type is not specified, fibre is used first to verify the connection, and if no connection is found, iSCSI is used.

EXAMPLE

The example below uses the LUN's World Wide Name to verify that the LUN is reachable, and logs any events to the SP event log. If the LUN is reachable, the software displays the following message: This example verifies two LUNs using the SP name and LUN number. If one or both LUNs are not reachable, the software displays a message similar to the following:

```
Storage device is reachable.  
Error: sancopy command failed  
Error returned from Agent SP A: CPM unable to locate device. Check  
that the  
device with this WWN exists (0x712ac007)
```

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h 11.12.13.14 sancopy -verify -wwn  
60:06:01:FA:2C:09:00:00:2B:89:4C:E9:65:D7:11 -log
```

```
naviseccli -h 11.12.13.14 sancopy -verify -lunnumber 12.34.56.78 6  
12.34.56.78 12
```


CHAPTER 12

EMC SnapView Software

This chapter introduces the EMC® SnapView™ software and its user interface, as well as the two command line interfaces for it. The command line interfaces include the server-based admsnap utility and the EMC VNX CLI for Block interface.

Note

Unless indicated otherwise, references to LUNs include classic LUNs and thin LUNs. If your system includes FAST Cache (solid state disk drives with flash memory, or SSD drives), refer to the Best Practices documentation for information on using these disks with SnapView, which is available on EMC Online Support.

Major topics are:

- [SnapView clone CLI commands](#)..... 526
- [Admsnap commands for clones](#)..... 554
- [SnapView Snapshot CLI commands](#).....560
- [Admsnap commands for snapshots](#)..... 583

SnapView clone CLI commands

This chapter explains the Unisphere SnapView clone CLI commands that manage SnapView clones.

Setting up SnapView to use clones

This section describes how to set up SnapView to use clones.

Note

You cannot expand the capacity of a LUN that is participating in a clone group until you remove the clone from the clone group and destroy the clone group. Neither the production nor secondary server can access this added capacity until the expansion is complete and you perform some additional operations. For detailed information on expanding a LUN, see [vnx_bcli_c_cli_cmd_ref_navi_cli_chap on page 23](#).

Prerequisites for setting up clones

Before you can set up and use clones, the following must be true:

- Source LUNs must be bound. For a client or production server to access a source LUN, you must assign the source LUN to a storage group and connect the storage group to the production server. To do this, you must enable data access control on the storage system.
- LUNs that you plan to use as clone LUNs must be bound. These LUNs must be the same size as the source LUNs that you will clone. EMC strongly recommends that you bind your clone LUNs on storage pools (RAID groups or thin pools) that are different from their source LUNs. The clone's storage pool does not have to be the same RAID type as the source LUN.
- For VMware ESX Servers, verify that the source LUN is presented to the virtual machine (guest operating system running on the virtual machine). For information on how to present a LUN to the virtual machine, refer to the VMware documentation that shipped with your ESX Server.
- For configuration restrictions when using VMFS volumes, go to the EMC E-Lab™ Interoperability Navigator on EMC Online Support and under the PDFs and Guides tab, open the VMware ESX server.pdf file.
- For a secondary server to access the clone LUN, the clone must be assigned to a storage group (but you cannot read the clone until you fracture it). The storage group must be connected to the secondary server that will access the clone. You must assign the clone LUN to a storage group other than the storage group that holds the source LUN. EMC supports placing a clone in the same storage group as its source LUN only if you use Replication Manager or Replication Manager/SE to put the clone in the storage group. This software provides the same host access to the clone and the source LUN. For information on using these software products, refer to the documentation for the product.

If you have a VMware ESX Server, the clone and source LUNs must be accessed by different virtual machines, unless the virtual machine is running one of the software programs that support same host access.

Overview of setting up SnapView to use clones

The following is a checklist for setting up SnapView clones:

- Allocate clone private LUNs; see [clone - snapview -allocatecpl on page 531](#).

Note

This step is required only in the initial setup of clones.

- Create a clone group; see [clone - snapview -createclonergroup on page 537](#).
- Add a clone to a clone group; see [clone - snapview -addclone on page 528](#).

CLI clone command functions

The navisecli clone functions, listed alphabetically, are:

Note

As of version 6.26, you can issue clone commands using the clone or snapview prefix.

- clone | snapview -addclone
- clone | snapview -allocatecpl
- clone | snapview -changeclone
- clone | snapview -changeclonefeature
- clone | snapview -changeclonergroup
- clone | snapview -consistentfractureclones
- clone | snapview -createclonergroup
- clone | snapview -deallocatecpl
- clone | snapview -destroyclonergroup
- clone | snapview -fractureclone
- clone | snapview -listclone
- clone | snapview -listclonefeature
- clone | snapview -listclonergroup
- clone | snapview -removeclone
- clone | snapview -reversesyncclone
- clone | snapview -syncclone

The CLI clone functions are listed functionally, in the order you would perform them, in the following table.

Table 49 navisecli SnapView clone command functions, listed in order of use

Essential functions (in order performed)	
clone snapview -allocatecpl	Designates clone private LUNs.
clone snapview -createclonergroup	Creates a clone group.
clone snapview -addclone	Adds a clone LUN to the clone group.
clone snapview -syncclone	Starts synchronizing the clone with the source LUN.
clone snapview -consistentfractureclones or clone snapview -fractureclone	Frees the clone LUN(s) from its source, allowing you to use the clone independently.

Table 49 navisecli SnapView clone command functions, listed in order of use (continued)

clone snapview -removeclone	Deletes a clone from a clone group.
clone snapview -destroyclonegroup	Deletes a clone group.
Optional and status functions (alphabetically)	
clone snapview -listclone	Displays clone information.
clone snapview -listclonefeature	Displays clone features such as whether creating a clone group on this system is supported and other useful information.
clone snapview -listclonegroup	Displays clone group information.
Optional reconfiguration functions (alphabetically)	
clone snapview -changeclone	Changes one or more properties of a clone.
clone snapview -changeclonefeature	For a system, enables or disables clone features such as protected restore.
clone snapview -changeclonegroup	Changes clone group settings.
clone snapview -deallocatecpl	Frees LUNs that were allocated as clone private LUNs from the system clone structure.
clone snapview -reversesyncclone	Starts a reverse sync process, in which the software copies the clone data to the source.

clone | snapview -addclone

Adds a clone to a clone group.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `navisecli clone` or the `navisecli snapview` command with `-addclone` adds a clone to a clone group. Each clone LUN you add must be the same size as the source LUN. When you add a clone to a clone group, you create an actual copy of the source LUN you specified when creating the group. The copy of the source LUN begins when you specify the `-issyncrequired 1` command when adding the clone LUN. This command is also used to add an initializing thin LUN as a clone target, provided the bundle containing the thin LUN support is committed in all the systems.

The software assigns each clone a clone ID. This ID remains with the clone until you remove the clone from its group.

Note

When you add a clone to a clone group, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

Source LUN writes and unfractured clone states — While the clone is part of the clone group and unfractured (not accessible to a secondary server), server writes to the source LUN can continue, with the following results:

Server writes to the source — When a source LUN receives a server write request, the clone transitions into a consistent state because the clone is no longer a byte-for-byte copy of its source. Writes made to the source LUN are simultaneously copied to the clone. The clone remains in a consistent state if server writes are made to the source LUN during the duration of the quiesce threshold. The quiesce threshold is the amount of time the clone must wait before transitioning from a consistent state to a synchronized state. If the source LUN receives any server writes during this period of time, the time set for the quiesce threshold restarts. You specify the quiesce threshold when you create a clone group.

Note

When you add a clone to the clone group, with the Initial Sync Required property selected, the clone state is synchronizing. The software transitions the clone to synchronized or consistent state only after the initial synchronization is complete.

No server writes to the source — If a source LUN does not receive any server writes for the duration of the quiesce threshold, any unfractured clone in a consistent state transitions into a synchronized state.

Note

See Clone states in *EMC VNX for Block Concepts* for information about the state of a clone.

SYNTAX

```
clone | snapview -addclone -name name|-CloneGroupUid uid -luns
lun_numbers [-issyncrequired 0|1]
[-RecoveryPolicy manual|auto]
[-SyncRate high|medium|low|value]
[-UseProtectedRestore 0|1]
```

OPTIONS

-name *name*

Specifies the name that you gave to the clone group with the `-createclonegroup` function. Use either this name or the clone group unique ID (16-byte WWN).

-CloneGroupUid *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID (refer to [clone | snapview -listclonegroup on page 546](#) for more information).

-luns *lun_numbers*

Specifies the LUN numbers that can serve as clones. Currently only one LUN is supported per command.

-issyncrequired 0|1

Specifies whether the software must synchronize the newly added clone.

0 - Synchronization is not required. Using this value lets you skip the synchronization period if the source LUN contains null data (as with a newly bound source LUN).

⚠ CAUTION

When you specify `issyncrequired 0`, the software marks the clone as synchronized regardless of its contents. Therefore, do not specify `issyncrequired 0` if the source LUN contains any data, because the clone will appear as synchronized when it is not synchronized, and will probably not contain usable data.

1 - Synchronization is required (default). The software starts synchronizing the clone immediately.

`-RecoveryPolicy manual|auto`

Specifies the synchronization action the system will take after a hardware component (such as an SP) fails. Values are:

- manual - Manual recovery. The administrator must act to start synchronization.
- auto - Automatic. Synchronization will start automatically using the other path (if available). This is the default.

The value you select also applies if reverse-synchronization was in progress before the failure.

`-SyncRate high|medium|low|value`

Sets the interval between writes to a clone being synchronized to high, medium, or low, or a value in milliseconds. If the value is larger than 100, the interval is in milliseconds. Otherwise, it is the scaled between high rate and medium rate - the value 1 is equivalent to high rate and 100 to medium rate. Default is medium. A high rate synchronizes the clone the fastest, but may affect performance on other system resources. A low value has minimal effect on performance, but prolongs the synchronization process.

Note

EMC recommends that you do not use a high synchronization rate on a system with a single SP.

`-UseProtectedRestore 0|1`

Determines the way the software reacts when I/O is received during a reverse sync operation.

- 0 - Updates the clone with the incoming server I/O. This is the default.
- 1 - Does not update the clone with the I/O.

EXAMPLE

This example adds the current logged-in user to the security file. Once you enter this command, you do not have to re-enter your username, password, or scope in the command line.

```
naviseccli -password password -scope 0 -addusersecurity
```

```
naviseccli -address ssl_spa clone -allocatecpl -spA 54 -spB 55
```

```
naviseccli -address ssl_spa clone -createclonegroup -name db3_clone  
-luns 22
```

```
naviseccli -address ssl_spa clone -addclone -name db3_clone -luns 43  
-o
```

```
naviseccli -address ssl_spa clone -listclonegroup
```

```
Name: db3_clone
CloneGroupUid: 50:06:01:60:20:04:21:3E:0C:00:00:00:00:00:00
InSync: No
```

These commands allocate clone private LUNs (LUNs 54 and 55), create a clone group named `db3_clone` for source LUN 22, add a clone LUN (LUN 43) to the group, and then list clone groups. Synchronization starts by default with the `-addclone` command since that command does not specify `-issyncrequired 0`. To remove the logged in user from the security file, issue the following last command:

```
naviseccli -removeusersecurity
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -allocatecpl

Reserves a clone private LUN for use with clones.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

Each clone private LUN must be a minimum of 1 gigabyte. Clone private LUNs larger than 1 gigabyte provide no performance benefit. You must bind these LUNs before you allocate them as a clone private LUN. You can use any LUN that is at least 1 gigabyte in size as a clone private LUN, except for the following:

- Thin LUNs
- Hot spare LUNs
- MirrorView remote mirror LUNs (LUNs used as either a primary or secondary image)
- SnapView clone, snapshot, or source LUNs
- SAN Copy source or destination logical units
- Private LUNs (LUNs reserved as clone private LUNs or for use by the reserved LUN pool)

You should bind clone private LUNs in a RAID group that normally does not see heavy I/O.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-allocatecpl` allocates or reallocates a clone private LUN on both SPs for clone operations. You must allocate one clone private LUN for each SP before you can create a clone group. This command supports thin LUNs, provided the bundle containing the thin support is committed in all the systems.

Note

The `clone -allocatecpl` command restricts Clone Private LUNs to 128 element size. If you select LUN(s) with element size not equal to 128, the system will display the following error message:

```
Invalid element size for chosen LUN(s). Please choose LUN(s) with
element size of 128.
```

If you use `-allocatecpl` to reallocate (change an existing allocation) clone private LUN, you must change the clone private LUN for both SPs. When you reallocate the clone

private LUNs, the software transfers the information stored in the previous clone private LUNs to the new clone private LUNs.

Note

When you allocate clone private LUNs, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
clone | snapview -allocatecpl -spA lun_number -spB lun_number
-o
```

OPTIONS

-spA lun_number

lun_number identifies the LUN to be used for the clone private LUN for SP A. The LUN must be bound. You must specify a clone private LUN for both SP A and SP B. However, the ownership of these LUNs does not matter, since the software will trespass each LUN you specify in each *-sp* switch to the specified SP.

-spB lun_number

lun_number identifies the LUN to be used for the clone private LUN for SP B. The LUN must be bound. See comments under *-spA*.

-o

Executes the command without prompting for confirmation.

EXAMPLE

This example allocates a clone private LUN to SP A (the LUN with LUN ID 54) and SP B (LUN with ID 55) for clone operations (see also the *-addclone* example):

```
naviseccli -address ssl_spa clone -allocatecpl -spA 54 -spB 55 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -changeclone

Changes one or more properties of a clone. The `clone -changeclone` or `snapview -changeclone` command is used with `naviseccli` as follows.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-changeclone` changes one or more clone properties. The clone can be in any state when you use this function. If you omit the override option (`-o`), the software displays a message to confirm the change.

SYNTAX

```
clone | snapview -changeclone -name name|-CloneGroupUid uid
-cloneid id [-RecoveryPolicy manual|auto]
[-SyncRate high|medium|low|value]
[-UseProtectedRestore 0|1]
[-o]
```

OPTIONS**-name** *name*

Specifies the name that you gave to the clone group with the `-createclonegroup` function. Use either this name or the clone group unique ID (16-byte WWN).

-CloneGroupUid *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID.

-cloneid *id*

Specifies the 16-character clone ID the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

-RecoveryPolicy *manual|auto*

Specifies the synchronization action the system will take after a hardware component (such as an SP) fails. Values are:

- `manual` - Manual recovery. The administrator must act to start synchronization.
- `auto` - Automatic. Synchronization will start automatically using the other path (if available). This is the default.

The value you select also applies if reverse-synchronization was in progress before the failure.

-SyncRate *high|medium|low|value*

Sets the interval between writes to a clone being synchronized to high, medium, or low, or a value in milliseconds. If the value is larger than 100, the interval is in milliseconds. Otherwise, it is the scaled between high rate and medium rate - the value 1 is equivalent to high rate and 100 to medium rate. Default is medium. A high rate synchronizes the clone fastest, but may degrade overall performance. A low value has minimal effect on performance, but prolongs the synchronization process.

-UseProtectedRestore *0|1*

Specifies the way the software will react when I/O is received during a reverse sync operation.

0 - Updates the clone with the I/O. Default.

1 - Does not update the clone with I/O.

This property is available only if the option is supported on the targeted system.

Note

When you change the property for the `-UseProtectedRestore` option, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

-o

Executes the command without prompting for confirmation.

EXAMPLE

This command changes the sync rate for the clone 0100000000000000 in clone group db3_clone.

```
naviseccli -address ssl_spa clone -changeclone -name db3_clone
-cloneid 0100000000000000 -syncrate high -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -changeclonefeature

Enables or disables clone features such as protected restore.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-changeclonefeature` can globally enable or disable the `-AllowProtectedRestore` option. This option allows you to use the protected restore feature on a per clone basis. For information on the protected restore feature, refer to [clone | snapview -reversesyncclone on page 550](#).

If you omit the override option (`-o`), the software displays a message to confirm the change.

SYNTAX

```
clone | snapview -changeclonefeature [-AllowProtectedRestore 0|1] [-o]
```

OPTIONS

`-AllowProtectedRestore 0|1`

Enables or disables the ability to use protected restores on LUNs in a system. Be aware that enabling protected restores reserves SP memory and prevents it from being used for any other purpose.

0-Does not allow protected restores on this system. Default.

1-Allows protected restores on this system.

Note

When you enable the `-AllowProtectedRestore` option, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

This example allows protected restores on the system that contains the SP `ss1_spa`.

```
naviseccli -address ss1_spa clone -changeclonefeature
-AllowProtectedRestore 1
```

```
Do you really want to perform this operation? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -changeclonegroup

Changes one or more parameters of an existing clone group.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-changeclonegroup` can change certain settings of an existing clone group.

SYNTAX

```
clone | snapview -changeclonegroup -name name|-CloneGroupUid
uid [-description description]
[-newname name]
[-o]
[-quiescethreshold seconds]
```

OPTIONS

`-name name`

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

`-CloneGroupUid uid`

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID.

`-description description`

Lets you enter a description about the clone group—up to 256 characters. If you want to include any spaces, enclose the entire description in double quotation marks. The quotation marks do not count towards the 256-character limit.

`-newname name`

Lets you specify a new name for the clone group. A valid name consists of at least one printable character and must not exceed 64 characters. Each clone group name must be unique on its system.

`-o`

Executes the command without prompting for confirmation.

`-quiescethreshold seconds`

Lets you specify the amount of time that must pass after an I/O to the source LUN before the software places any of its unfractured clones in the synchronized state. The minimum value is 10, the maximum value is 3600.

EXAMPLE

This example changes the quiesce threshold for clone group `db3_clone` to 30 seconds.

```
naviseccli -address ssl_spa clone -changeclonegroup -name db3_clone
-quiescethreshold 30
```

```
Do you really want to perform this operation? y
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -consistentfractureclones

Splits a set of clones from their source LUNs.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

The clones you want to fracture must be within different clone groups. You cannot perform a consistent fracture on clones belonging to different systems. You must specify two or more clones that you want to fracture or an error will be displayed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-consistentfractureclones` fractures more than one clone at the same time in order to preserve the point-in-time restartable copy across the set of clones. The SnapView driver will delay any I/O requests to the source LUNs of the selected clones until the fracture has completed on all clones (thus preserving the point-in-time restartable copy on the entire set of clones).

Note

A restartable copy is a data state having dependent write consistency and where all internal database/application control information is consistent with a database management system/application image.

After the consistent fracture completes, no group association exists between the clones. If a failure occurs on any of the clones, the consistent fracture will fail on all of the clones. If any clones within the group were fractured prior to the failure, the software will resynchronize those clones. For more information on the effects of fracturing a clone, refer to the `-fractureclone` description in [clone | snapview -fractureclone on page 540](#).

SYNTAX

```
clone | snapview -consistentfractureclones
-CloneGroupNameCloneId name cloneId|-CloneGroupUidCloneId CGUId cloneId [-o]
```

OPTIONS

`-CloneGroupNameCloneId name cloneId`

Lets you specify the clone group name and the clone ID for the clones you want to fracture. If you specify the clone group name and clone ID for one clone, you must specify the same information for all the other clones you want to fracture at the same time.

The clone group name is the name you specified when you created the clone group. The clone ID is the 16-character ID. The system assigned when the clone was added to the clone group. The clone ID is always numeric.

Use the `-listclonegroup` command to display all the clone group names and clone IDs on the system (see [clone | snapview -listclonegroup on page 546](#)).

`-CloneGroupUidCloneId CGUId cloneId`

Lets you specify the clone group ID and the clone ID for the clones you want to fracture. If you specify the clone group ID and clone ID for one clone, you must specify the same information for all the other clones you want to fracture at the same time.

The clone group ID is the 16-byte unique ID the software assigned when you created the clone group. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). The clone ID is the 16-character ID the system assigned when the clone was added to the clone group. The clone ID is always numeric.

Use the `-listclonegroup` command to display all the clone group IDs and clone IDs on the system (see [clone | snapview -listclonegroup on page 546](#)).

-o

Executes the command without prompting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -consistentfractureclones
-CloneGroupNameCloneId CGname1 0100000000000000 CGname2
0200000000000000 CGname3 0300000000000000 -o
```

```
naviseccli -address ssl_spa snapview -consistentfractureclones
-CloneGroupUidCloneId 50:06:01:60:80:60:03:E5:3A:00:00:00:00:00:00:00
0100000000000000 50:06:01:60:80:60:03:E5:3A:00:00:00:00:00:00:00
0200000000000000 50:06:01:60:80:60:03:E5:3A:00:00:00:00:00:00:00
0300000000000000 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -createclonegroup

Creates a new clone group based on a source LUN.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

You must have allocated a clone private LUN for each SP.

Any source LUN that is at least 1 GB in size is eligible to be cloned, except for the following:

- Hot spare LUNs
- Clone LUNs (LUNs participating in any clone group as either a source LUN or a clone LUN)
- Snapshot LUNs
- MetaLUNs
- Private LUNs (LUNs reserved as clone private LUNs, in a reserved LUN pool, or in a write intent log)
- A LUN in the process of a migration

Note

With SnapView version 6.24 or later, a clone's source LUN can be a remote mirror LUN (a LUN participating as either a primary or secondary image).

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-createclonegroup` creates a new clone group using an existing LUN as a source.

Note

When you create a clone group, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
clone | snapview -createclonegroup -name name -luns lun_number
[-description description]
[-quiescethreshold seconds] [-o]
```

OPTIONS

-name *name*

Lets you specify the name of the clone group. A valid name consists of at least one printable character and must not exceed 64 characters. Each clone group name must be unique on its system.

-luns *lun_number*

The *lun_number* identifies the source LUN of the clone group.

-description *description*

Lets you enter a description about the clone group—up to 256 characters. If you want to include any spaces, enclose all text with double quotation marks.

-quiescethreshold *seconds*

Lets you specify the amount of time that must pass after an I/O to the source LUN before the software places any of its unfractured clones in the synchronized state. The minimum value is 10, the maximum value is 3600.

-o

Executes the command without prompting for confirmation.

EXAMPLE

This command creates a clone group named `db3_clone` for source LUN 22. Refer to the example for `-addclone`, [clone | snapview -addclone on page 528](#).

```
naviseccli -address ssl_spa clone -createclonegroup -name db3_clone
-luns 22 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

POSTREQUISITES

You will need to add a clone using the `clone | snapview -addclone` command and synchronize the clone before you can fracture and use it.

You can use the `-changeclonegroup` function to change some of the create parameters.

clone | snapview -deallocatecpl

Deallocates the clone private LUNs.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

You must remove all clones and destroy any clone groups before deallocating the clone private LUNs. You must also globally disable the Allow Protected Restore option. To do this, issue the `clone | snapview -changeclonefeature -AllowProtectedRestore 0` command (refer to [clone | snapview -changeclonefeature on page 534](#) for more information on this command).

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-deallocatecpl` removes the LUNs that were previously allocated as clone private LUNs. You can also use the `-deallocatecpl` command to reallocate clone private LUNs, for instance, if a clone private LUN fails or if you decide that you want to use a different LUN as your clone private LUN. If you use the `-deallocatecpl` command to reallocate the clone private LUNs, you can do so with existing clone groups and clones. When you reallocate the clone private LUN, the software transfers the information stored in the previous clone private LUN to the new clone private LUN.

Note

If you reallocate one clone private LUN, you must reallocate the other clone private LUN.

SYNTAX

```
clone | snapview -deallocatecpl [-o]
```

OPTIONS

`-o`

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -deallocatecpl
-o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -destroyclonegroup

Deletes a clone group.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

Before you can destroy a clone group, you must remove all clones in the clone group.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-destroyclonegroup` permanently removes a clone group from the system and releases the source LUN specified when you created the clone group.

SYNTAX

```
clone | snapview -destroyclonegroup -name name|-CloneGroupUid uid
[-o]
```

OPTIONS

`-name name`

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

`-CloneGroupUid uid`

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-`

`listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup](#) on page 546).

-o

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -destroyclonegroup -name db3_clone
-o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -fractureclone

Splits a clone from its source LUN.

PREREQUISITES

You must have a user account on the system on which you want to execute the command.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-fractureclone` splits a clone from its source LUN, allowing a secondary server to access the clone as an independent LUN.

Note

To fracture more than one clone at the same time, you must use the `-consistentfractureclones` command (see [clone - snapview - consistentfractureclones](#) on page 536).

Fracturing a clone or a set of clones (`-consistentfracture`) does not remove the clone(s) from its clone group. It remains a part of the clone group — although other servers can access it as an independent LUN — until you explicitly remove it (`-removeclone`). While a clone remains in its clone group, you can synchronize it again using just the `-syncclone` function.

Note

Unless you have additional software that supports same host access, you must assign the clone LUN to a storage group other than the storage group that holds the source LUN(s). You also must assign multiple fractured clones, of the same source LUN(s), to different storage groups. See [Prerequisites for setting up clones](#) on page 526 for information on software that supports same host access.

Two types of fractures can occur - administrative or system.

An administrative fracture will occur after any of the following events:

- The production server performed a fracture operation on the clone or set of clones.
- The secondary server performed a reverse synchronization operation on another clone in the clone group.
- The clone driver receives an error reading or writing to the source or clone during a synchronization or reverse synchronization. Possible errors include a failed sector or a bad block.

Note

To verify if an error occurred, issue the `-listclonegroup` command and determine if the value for `SourceMediaFailure` or `CloneMediaFailure` is set to `TRUE`. If an error did occur, you must correct the failure, then re-issue the synchronization or reverse synchronization operation. If the error persists, contact your EMC service provider.

A system fracture will occur after any of the following events:

- The clone driver encountered an internal driver error and fractured the clone as a result. If this occurs, you must correct the error condition, perform a fracture (administrative fracture), and then synchronize the clone.
- The source LUN trespasses to the peer SP, either manually or due to an SP failure. If this occurs, the system fracture is temporary and the clone will automatically synchronize with its source LUN when the source trespasses back to its SP.
- Other replication software uses the clone or its source, for example, to start a SAN Copy session or to create a snapshot of the LUN. If this occurs, the system fracture is temporary and the clone will automatically synchronize with its source LUN after the event completes.

You can fracture a clone in any state. However, to make a fractured clone available for server write requests, you must fracture the clone while it is in a synchronized or consistent state. Once you fracture the clone in either of these states, it transitions to a consistent state.

If you want the clone to contain a predictable point-in-time copy, you must ensure that no server writes are occurring when you fracture the clone. You can determine that no writes have been issued to the source LUN for a period of time equal to the quiesce threshold if you fracture the clone while it is in a synchronized state. However, it does not guarantee that another server write will not occur around the same time you fracture the clone.

If the clone never enters the synchronized state, server writes are continuing. Any server writes to the source that occur after the fracture are not copied to the clone. If you believe you have stopped all I/O to the clone's source LUN, it may be difficult to determine the reason for those writes. For example, they can be due to heartbeats in some clustered systems.

In some cases it is useful to fracture the clone in a consistent state. For instance, you would fracture a clone in a consistent state when:

- you can determine that there are no writes being issued to the source LUN and you prefer not to wait for the quiesce threshold to elapse.
- your application has a hot backup mode in which the application continues to write to the source, but internally understands how to disregard writes generated after your specified point in time.
- you need only a crash consistent image (the image you would have if the server failed or shut down improperly).

If the clone synchronization is detracting from the system's performance, you can fracture the clone in any other state and then synchronize it at a time when performance is not an issue.

Writes to the fractured clone LUN

When a fractured clone receives a server write request, the software marks the clone as dirty, which indicates that the clone has been modified and is no longer a byte-for-byte copy of its source LUN.

Writes to the source LUN

After you fracture the clone, the software does not copy any server write requests made to the source LUN unless you manually perform a synchronization. Synchronizing a

fractured clone unfractures the clone and updates the contents on the clone with its source LUN.

SYNTAX

```
clone | snapview -fractureclone -name name | -CloneGroupUid uid -cloneid id [-o]
```

OPTIONS

-name *name*

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

-CloneGroupUid *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).

-cloneid *id*

Specifies the 16-character clone ID the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

-o

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -fractureclone -name db3_clone -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -listclone

Displays information about a specific clone LUN.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-listclone` displays information about one clone, not clone groups. You must specify the clone name or clone group ID. You can obtain more specific information with function switches.

SYNTAX

```
clone | snapview -listclone -Name name | -CloneGroupUid uid
[-all]
[-AvailableForIO]
[-CloneCondition]
[-cloneid id]
[-CloneLuns]
[-CloneLunWWNs]
[-CloneState]
[-CloneMediaFailure]
[-IsDirty]
[-IsFractured]
```

```
[-PercentSynced]
[-RecoveryPolicy]
[-SyncRate]
[-timeoffracture]
[-UseProtectedRestore 0|1]
```

OPTIONS

- Name *name*
Specifies the name of the clone group. This is required for getting information on a clone.
 - CloneGroupUid *uid*
Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).
 - all
Lists all the properties of a clone without having to type in any other switch. This switch must be specified to know the `-timeoffracture` of the clone.
 - AvailableForIO
Displays whether the clone is available for I/O from an attached server (*yes* or *no*).
 - CloneCondition
Displays the clone condition. Typically, this shows the nature of a faulted clone. The values include Administratively Fractured, System Fractured, Waiting on administrator to initiate synchronization, Queued to be synchronized, Synchronizing, Normal, Unknown.
-
- Note**
- If you fractured a clone (`-fractureclone`) or a group of clones (`-consistentfractureclones`), the clone condition will be displayed as Administratively Fractured.
-
- cloneid *id*
Specifies the 16-character clone ID the system assigned when the clone was added. You can learn this using the `-listclonegroup` command function. The id is case sensitive.
 - CloneLuns
Displays a list of clone LUN numbers.
 - CloneLUNWWNs
Displays the WWN of the clone.
 - CloneState
States are consistent, out-of-sync, remote mirror synchronizing, reverse-out-of-sync, reverse-synchronizing, synchronized or synchronizing.
 - CloneMediaFailure
Indicates whether an error occurred during I/O with the clone LUN (*yes* or *no*).
 - IsDirty
Indicates whether the clone received a server write request, which means it is not a byte-for-byte duplicate of its source LUN (*yes* or *no*).
 - IsFractured
Indicates whether the clone was fractured (*yes* or *no*).
 - PercentSynced
Displays the percentage of the clone that is synchronized with its source.
 - RecoveryPolicy
Indicates the recovery policy (see `-addclone`), Automatic or Manual.

- SyncRate
Displays high, medium, low or a user-specified positive integer.
- timeoffracture
Displays the date/time when the clone images were administrative fractured.
- UseProtectedRestore 0|1
Indicates the state of the Protected Restore option. Values are 0 (clone may be updated with new I/O during a restore) or 1 (clone cannot be updated with new I/O during restore).

EXAMPLE #1

```
naviseccli -address ssl_spa clone -listclone -name db3_clone
```

```
Name: db3_clone
CloneGroupUid: 50:06:01:60:80:60:03:E5:3A:00:00:00:00:00:00:00
CloneID: 010000000000000000
CloneState: Synchronized
CloneCondition: Normal
AvailableForIO: No
CloneMediaFailure: No
IsDirty: No
IsFractured: No
PercentSynced: 100
RecoveryPolicy: Auto
SyncRate: Medium
CloneLUNs: 1
UseProtectedRestore: No
```

EXAMPLE #2

```
naviseccli -scope 0 snapview -listclone -name CG -timeoffracture
```

```
Name: CG
CloneGroupUid: 50:06:01:60:BC:E0:1C:43:01:00:00:00:00:00:00:00
CloneID: 020000000000000000
TimeofFracture: 12/15/09 11:25:19
```

EXAMPLE #3

```
naviseccli -scope 0 snapview -listclone -name CG -all
```

```
Name: CG
CloneGroupUid: 50:06:01:60:BC:E0:1C:43:01:00:00:00:00:00:00:00
CloneID: 020000000000000000
CloneState: Consistent
CloneCondition: Administratively Fractured
AvailableForIO: Yes
CloneMediaFailure: No
IsDirty: No
IsFractured: Yes
PercentSynced: N/A
RecoveryPolicy: Auto
SyncRate: Medium
CloneLUNs: 10
UseProtectedRestore: No
TimeofFracture: 12/15/09 11:25:19
```

clone | snapview -listclonefeature

Displays general information for a system that supports clones.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-listclonefeature` and no other switches lists general clone information for a system that supports clones. You can obtain more specific information with switches.

SYNTAX

```
clone | snapview -listclonefeature [-AllowProtectedRestore]
[-AvailableLUNs]
[-CanCreate]
[-Cpl]
[-CplSPA]
[-CplSPB]
[-DriverVersion]
[-MaxCloneGroupsUsingCPL]
[-MaxCloneMemory]
[-MaxClonesPerGroup]
[-MaxLUNsPerSource]
[-MinCplSize]
```

OPTIONS

- `-AllowProtectedRestore`
Displays the state of the system Protected Restore option. The state can be `no`, which indicates that the system should update the clone with the I/O, or `yes`, which indicates that it should not update the clone with the I/O. This option is available only if it is supported on the targeted system.
- `-AvailableLUNs`
Lists LUN numbers currently available to be used as either a source or clone.
- `-CanCreate`
Indicates if a clone group can be created on this system. Valid answers are `yes` or `no`.
- `-Cpl`
Indicates whether clone private LUNs are allocated and enabled on this system. Valid answers are `yes` or `no`.
- `-CplSPA`
Displays the LUN number of the clone private LUN for SP A.
- `-CplSPB`
Displays the LUN number of the clone private LUN for SP B.
- `-DriverVersion`
Displays the version of the clone driver.
- `-MaxCloneGroupsUsingCPL`
Displays the maximum number of clone groups that will use the clone private LUNs that this system supports.
- `-MaxCloneMemory`
Displays the maximum amount of memory that the clone driver can allocate.
- `-MaxClonesPerGroup`
Displays the maximum number of clones per clone group.
- `-MaxLUNsPerSource`
Displays the maximum number of LUNs allowed in a source.
- `-MinCplSize`
Displays the minimum size allowed, in blocks, for each clone private LUN.

EXAMPLE

```
naviseccli -address ssl_spa clone -listclonefeature
```

```
CanCreate: No
MaxCloneGroupsUsingCPL: 50
MaxClonesPerGroup : 8
```

```

MaxLUNsPerSource: 1
MaxCloneMemory: 10485760
Cpl: Yes
CplSPA: 2
CplSPB: 3
MinCplSize: 250000
DriverVersion: 2
AvailableLUNs:
AllowProtectedRestore: No

```

clone | snapview -listclonegroup

Displays clone group information.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-listclonegroup` and no other switches lists general information for a clone group. You can obtain more specific information with switches.

SYNTAX

```

clone | snapview -listclonegroup [-all]
[-CloneCount]
[-CloneGroupUid]
[-Clones]
[-CloneLunWWNs]
[-Description]
[-InSync]
[-IsControllingSP]
[-Name [name]]
[-QuiesceThreshold]
[-SourceLUNSize]
[-SourceMediaFailure]
[-Sources]

```

OPTIONS

`-all`

Lists all the properties of a clone without having to type in any other switch. This switch must be specified to know the `-timeoffracture` of the clone.

`-CloneCount`

Displays the number of clones and clone names defined in the clone group.

`-CloneGroupUid`

Displays the Uid and clone name. The Uid is the 16-byte unique ID assigned when the clone group was created.

`-Clones`

Displays the following information about each clone:

Table 50 Clones information

CloneID	Identification number of the clone. This number is unique per clone group, not system.
CloneState	Values include consistent, out-of-sync, remote mirror synchronizing, reverse-out-of-sync, reverse-synchronizing, synchronized or synchronizing

Table 50 Clones information (continued)

CloneCondition	Values include Administratively Fractured, System Fractured, Waiting on administrator to initiate synchronization, Queued to be synchronized, Synchronizing, Normal, or Unknown.
AvailableForIO	Clone LUN availability for I/O from an attached server: <i>yes</i> or <i>no</i> .
CloneMedia Failure	Whether clone media failure occurred: <i>yes</i> or <i>no</i> .
IsDirty	Clone containing writes occurred from a source other than the clone group source: <i>yes</i> or <i>no</i> .
IsFractured	Clone fractured state: <i>yes</i> or <i>no</i> .
PercentSynced	Percentage of the clone instance that is currently synchronized with its source. Also applies to reverse-synchronizing operations.
RecoveryPolicy	A user-defined value: <i>Automatic</i> or <i>Manual</i> .
SyncRate	A user-defined value: <i>high/medium/low</i> .
CloneLUNs	List of clone LUN numbers.
UseProtected Restore	State of the clone Protected Restore option: <i>yes</i> or <i>no</i> .

-CloneLUNWWNs

Displays the WWN of the clone.

-Description

Displays the descriptions and clone names in the clone group.

-InSync

Indicates whether the source LUN is available for I/O.

-IsControllingSP

Indicates whether the SP specified in the command is the SP that owns the clone group (*yes* or *no*).

-Name *name*

Displays information for the clone group name as given to the clone group with the `createclonegroup` function.

-QuiesceThreshold

Displays the amount of time that must pass after an I/O to the source LUN before the software places any of its unfractured clones in the synchronized state. The minimum value is 10, the maximum value is 3600.

-SourceLUNSize

Displays the names and combined size of the source LUNs.

-SourceMediaFailure

Indicates whether an error occurred during I/O with the source LUNs.

-Sources

Displays the LUN numbers that constitute the source in the clone group.

EXAMPLE

```
naviseccli -address ssl_spa clone -listclonegroup -Name clone_0 -all
```

```
Name: clone_0
CloneGroupUid: 50:06:01:60:BB:20:10:DD:01:00:00:00:00:00:00:00
InSync: Yes
Description:
QuiesceThreshold: 60
SourceMediaFailure: No
```

```

IsControllingSP: No
SourceLUNSize: 2097152
CloneCount: 1
Sources: 4
Clones:
CloneID: 0100000000000000
CloneState: Consistent
CloneCondition: Administratively Fractured
AvailableForIO: Yes
CloneMediaFailure: No
IsDirty: No
PercentSynced: N/A
RecoveryPolicy: Auto
SyncRate: Medium
CloneLUNs: 5
UseProtectedRestore: No
IsFractured: Yes
TimeofFracture: 09/22/09 02:55:55

```

clone | snapview -removeclone

Removes a clone LUN from a clone group.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

To remove a clone from its clone group, the clone:

- must be fractured or in a synchronized or consistent state.
- cannot be in a synchronizing or reverse synchronizing state.
- cannot be in a queue to be synchronized.

Note

In an effort to control resource consumption, the SnapView driver limits a maximum of 40 concurrent synchronizations and/or reverse synchronizations per SP. Once this limit is met, the software queues all other synchronization and reverse synchronization requests.

- cannot be removed if it is in a reverse-out-of-sync state and it is the only clone in the clone group. If you remove a clone under these conditions, you cannot restore it. It is recommended that you perform a reverse synchronization operation from the clone and, when the reverse synchronization is successful, you should remove the last clone from the clone group.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-removeclone` removes a clone from its clone group. When you remove a clone from a clone group, the clone becomes a conventional (regular) LUN and is no longer associated with its source or clone group.

The clone group itself persists until you remove it using `-destroyclonegroup`.

SYNTAX

```
clone | snapview -removeclone -namename|-CloneGroupUIduid
-cloneid id [-o]
```

OPTIONS

`-name` *name*

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

`-CloneGroupUid` *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).

`-cloneid` *id*

Specifies the 16-character clone ID of the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

`-o`

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -removeclone -name db3_clone -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -resetfracturelog

Marks all data chunks of the source or clone LUN as being modified in the clone private LUN log. This allows you to perform a full synchronization or reverse-synchronization.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

You can use this switch only if the clone was administratively fractured and the clone is in an out-of-sync or reverse-out-of-sync state.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-resetfracturelog` simulates server writes to the source or clone LUN, marking all data chunks on the LUN as being modified (changed).

SYNTAX

```
clone | snapview -resetfracturelog -name name | -CloneGroupUid uid  
-cloneid id [-o]
```

OPTIONS

`-name` *name*

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

`-CloneGroupUid` *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).

`-cloneid` *id*

Specifies the 16-character clone ID the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

-o

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

```
naviseccli -address ssl_spa clone -resetfracturelog -name db3_clone
```

OUTPUT

None if the command succeeds; status or error information if it fails.

POSTREQUISITES

After issuing the `-resetfracturelog` command, you must issue a synchronization or reverse-synchronization. This will initiate a full synchronization or reverse-synchronization since all data chunks were marked as being modified.

Note

Performing multiple full synchronizations or reverse-synchronizations at the same time may affect performance on other system resources.

clone | snapview -reversesyncclone

Overwrites the source LUN with data from the clone LUN.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

NOTICE

You cannot reverse synchronize the clone if the clone's source is a MirrorView secondary image or a MirrorView primary image with unfractured secondaries.

To reverse synchronize a fractured clone, the following must be true:

- The clone is not in an out-of-sync state.
- Another clone in the clone group is not in a synchronizing or reverse synchronizing state.

For a Windows operating system, you must disable the indexing service and recycle bin on the source LUN of the clone. This prevents data corruption during reverse synchronization.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-reversesyncclone` replaces the data on the source LUN with the data on the clone. This allows you to revert to an earlier copy of the source LUN, for instance if the source became corrupted.

To ensure that there is no data corruption on the source LUN, you have to take the source LUN offline before you initiate the reverse synchronization. Once the operation begins, you can bring the source LUN back online.

When you initiate a reverse synchronization from Unisphere, the software immediately fractures all the clones in the clone group and then the software unfractures the clone from which you initiated the reverse synchronization. The software then copies the data chunks that differ between the source and the clone to the source LUN. The source LUN can instantly access the data, while the actual copying continues in the background.

Writes to the source LUN

Any server writes made to the source LUN after the clone was originally fractured are overwritten with the data on the clone. Any unmodified (unchanged) chunks of data are

not rewritten to the source. After the reverse synchronization has completed, the clone that initiated the reverse synchronization remains unfractured.

Note

If you issue the `-UseProtectedRestore 1` switch with the `clone | snapview -reversesyncclone` command, after the reverse synchronization has completed, SnapView fractures the clone that initiated the reverse synchronization.

During a reverse synchronization, the software automatically copies any server writes made from the source LUN to the clone. If you do not want source writes copied to the clone during a reverse synchronization, you must issue the `-UseProtectedRestore 1` switch with the `clone | snapview -reversesyncclone` command.

Note

If you modify the same data chunk on the clone more than once, the software copies only the last modification to the source LUN.

The Protected Restore feature protects the data on a clone during a reverse synchronization. When you select this feature, during a reverse synchronization the software will not copy any server writes made from the source LUN to the clone. Instead, the software records information in the clone private LUN to identify the source LUN writes for subsequent synchronizations.

Once you initiate a reverse synchronization, the software immediately unfractures the clone that initiated the reverse synchronization. Then the software fractures any other clones in the clone group in order to protect them from corruption should the reverse synchronization operation fail. The software then begins to copy its data to its source LUN. After the reverse synchronization has completed, the software fractures the clone that initiated the reverse synchronization.

You can enable the Protected Restore feature on a per-clone basis (and not on a per-clone-group basis). You can select this feature when you first add a clone to a clone group (`clone | snapview -addclone -UseProtectedRestore 1`) or at any time before you perform a reverse synchronization (`clone | snapview -listclone -UseProtectedRestore 1`).

Note

Before you can initiate the Protected Restore feature, you must globally enable it by issuing the `clone | snapview -changeclonefeature [-AllowProtectedRestore 1]` command. When you select this option, the SnapView driver automatically allocates 8 MB of additional memory per SP. The additional memory is fixed and is used to monitor modified blocks on the source LUN, in order to prevent these blocks from being overwritten by the clone during a reverse synchronization. This additional memory counts against the total replication software memory budget.

While the clone is reverse synchronizing, you cannot:

- add a clone to the clone group
- remove the clone that is reverse synchronizing from the clone group
- synchronize any clone in the clone group
- reverse synchronize any clone in the clone group

SYNTAX

```
clone | snapview -reversesyncclone -name name|-clonegroupUid
uid -cloneid id [-UseProtectedRestore 0|1]
```

[-o]

OPTIONS**-clonegroupUid** *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).

-cloneid *id*

Specifies the 16-character clone ID the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

-name *name*

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

-o

Tells the CLI to execute the command without waiting for confirmation.

-UseProtectedRestore 0|1

Specifies the way the software will react when I/O is received during a reverse synchronizing operation.

0-Updates the clone with I/O to the source LUN.

1-Does not update the clone with I/O to the source LUN.

This property is available only if the option is supported on the targeted system.

EXAMPLE

```
naviseccli -address ssl_spa clone -reversesyncclone -name db3_clone
-cloneid 0100000000000000 -o
```

OUTPUT

None if the command succeeds; status or error information if it fails.

clone | snapview -syncclone

Starts copying the data on a source LUN to its clone.

PREREQUISITES

You must have a user account on the system on which you want the command executed.

To synchronize a fractured clone, the following must be true:

- The clone cannot be in a reverse-out-of-sync state.
- Another clone in the clone group cannot be in a reverse synchronizing state.
- If a clone is system fractured, you must manually fracture the clone before you can synchronize it.

For a Windows operating system, the indexing service and recycle bin must be disabled on the clone LUN. This prevents data corruption during the synchronization.

DESCRIPTION

The `naviseccli clone` or the `naviseccli snapview` command with `-syncclone` begins the process of synchronization: the copying of data on the source LUN to its clone. The purpose of synchronizing a fractured clone is to update the contents

on the clone with its source LUN. Once you synchronize a fractured clone, the clone becomes unfractured and behaves the same as a newly added clone in a clone group.

When you initiate a synchronization, the software copies the data chunks that are different between the source LUN and the clone LUN to the clone. The software also copies any server writes from the source LUN to the clone LUN. The software does not copy any unmodified chunks to the clone.

Note

If the same chunk of data is modified on the source LUN more than once, only the last modification is copied to the clone.

While the clone is synchronizing, you:

- cannot remove the clone in a synchronizing state.
- cannot perform a reverse synchronization with any other clone in the clone group.

Note

In an effort to control resource consumption, the SnapView driver limits a maximum of 40 concurrent synchronizations and/or reverse synchronizations per SP. Once this limit is met, the software queues all other synchronization and reverse synchronization requests.

SYNTAX

```
clone | snapview -syncclone -name name|-CloneGroupUid uid
-cloneid id [-o]
```

OPTIONS

-name *name*

Identifies the clone group. Use the name assigned when the clone group was created with `-createclonegroup`. To identify a clone group, you can use either the clone name or the clone group unique ID.

-CloneGroupUid *uid*

Specifies the 16-byte unique ID the software assigned when the clone group was created. The clone group ID consists of the numbers 0 through 9 and the letters A through F (these letters are not case sensitive). You can use the `-listclonegroup` command function to obtain the clone group ID (see [clone | snapview -listclonegroup on page 546](#)).

-cloneid *id*

Specifies the 16-character clone ID the system assigned when the clone was added. You can obtain this using the `-listclonegroup` command function. The id is numeric.

-o

Tells the CLI to execute the command without waiting for confirmation.

EXAMPLE

Note

Before issuing the synchronization command, you should flush all cached data to the source LUN.

These commands list the state of all clone groups on system SP `ss1_spa`, and then start synchronization of the clone `db3_clone`.

```
naviseccli -address ss1_spa clone -listclonegroup
```

```
naviseccli -address ssl_spa clone -syncclone -name db3_clone -cloneid
0100000000000000 -o
```

```
Name: db3_clone
CloneGroupUuid: 50:06:01:60:20:04:21:3E:0C:00:00:00:00:00:00:00
InSync: Yes
Description:
QuiesceThreshold: 60
SourceLUNSize: 2097152
CloneCount: 2
Sources: 16
Clones:
CloneID: 0100000000000000
```

OUTPUT

None if the command succeeds; status or error information if it fails.

Admsnap commands for clones

This section explains how to use the admsnap utility software with SnapView clones.

About admsnap software and clones

The admsnap utility is an executable program that you can run interactively with a script to manage clones. The admsnap utility resides on the servers connected to the system with the SnapView driver. You must execute the following admsnap commands from one of these attached servers:

- `admsnap clone_activate` – scans for new devices and, for Windows, assigns a drive letter.
- `admsnap clone_deactivate` – flushes server buffers and removes the drive letter assigned with `clone_activate` (Windows only).
- `admsnap flush` – flushes data from a source LUN to its clone, so that the source and clone are in a synchronized state.
- `admsnap list` – lists devices and their World Wide Name (WWN) within the operating system.

Note

The admsnap software supports native device names and PowerPath pseudo-device names. If you use a PowerPath pseudo-device name, you can add this device name to the `etc/fstab` file (for Linux) or to the `etc/vfstab` file (for Solaris). Admsnap supports both disk-based naming scheme and enclosure-based naming scheme for Veritas devices.

For a step-by-step overview and examples on setting up and using clones with admsnap and VNX CLI for Block, refer to SnapView examples in *EMC VNX for Block Concepts*.

Requirements for admsnap and clones

Before you can use admsnap for clones, the following must be true.

General requirements

- The server must have access to VNX CLI for Block. This is required to create a prerequisite LUN called a clone private LUN and to create clones.

- The system must have the SnapView software installed.
- The production and secondary servers must have both the Unisphere Host Agent and admsnap software installed. To install the host agent and admsnap, refer to the EMC Block server support products guide for your operating system.

Platform-specific requirements

- For HP-UX, admsnap requires a volume group entry for each path to the system. The entry or entries can be contained in a single volume group or multiple volume groups. On the destination system, the volume group entry or entries must specify a clone.

admsnap clone_activate

Scans for new LUNs on the secondary server.

PREREQUISITES

You must have administrative privileges on the secondary server.

Before this command will work, a clone group must exist (`clone | snapview -createclonegroup`) and a clone must already exist (that is, you must start one by adding a LUN to the clone group using `clone | snapview -addclone`). Then you must fracture the clone and add it to a storage group that is connected to the secondary server.

DESCRIPTION

The `admsnap clone_activate` command scans for new devices and, for Windows, assigns a drive letter to every new device it finds. When you activate a clone and add it to a storage group, the clone becomes accessible on the secondary server.

Note

Although the `clone_activate` command is prefixed with the term `clone`, you may issue this command on any device that requires a scan for new LUNs and, for Windows only, for the assignment of a drive letter.

Depending on your operating system, the `admsnap clone_activate` command will do the following:

Table 51 Function of `admsnap clone_activate` command for an operating system

Server OS	Function description
Windows	Scans for new devices and assigns a drive letter to every new device it finds. For Windows Remote Desktop and Terminal Services - If a terminal server is running on the secondary server, you must remove and disable it before using the <code>clone_activate</code> command. If you have the terminal server enabled, it prevents the command from working properly.
AIX, HP-UX, and Solaris	Rescans the bus for new devices.
Linux	The <code>clone_activate</code> command is not supported. You must either reboot the server or unload and load the HBA driver.

Table 51 Function of admsnap clone_activate command for an operating system (continued)

Server OS	Function description
NetWare	The clone_activate command is not supported. Run the command list_devices or use the command scan_all_LUNs on the console.

SYNTAX

```
admsnap clone_activate [-h]
```

OPTIONS

-h

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. Admsnap will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, and AIX systems only.

EXAMPLE

```
admsnap clone_activate
```

admsnap clone_deactivate

Flushes server buffers and removes assigned drive letters. This command is supported on Windows systems only.

PREREQUISITES

You must have administrative privileges on the secondary server.

You must use this command on a secondary server after the clone has been activated and before you synchronize the fractured clone, or on the production server before you reverse synchronize a fractured clone.

DESCRIPTION

On a Windows server, the admsnap clone_deactivate command flushes all server buffers and removes the drive letter assigned by clone_activate. When you deactivate a clone, the software takes the clone offline making it inaccessible to the secondary server.

Note

Although the clone_deactivate command is prefixed with the term clone, you may issue this command on any Windows device that requires a volume to be taken offline, making it inaccessible to the secondary server. After you issue the deactivate command, you can remove the clone from the storage group.

Table 52 Function of admsnap clone_deactivate command for an operating system

Server OS	Function description
Windows	Flushes all server buffers and removes the drive letter assigned by clone_activate. This operation will attempt to take the device offline and make it inaccessible by the operating system. To bring the device back online, you must issue the clone_activate command.

Table 52 Function of `admsnap clone_deactivate` command for an operating system (continued)

Server OS	Function description
	<p>For Windows Remote Desktop and Terminal Services - If a terminal server is running on the secondary server, you must remove and disable it before using the <code>clone_deactivate</code> command. If you have the terminal server enabled, it prevents the <code>clone_deactivate</code> command from working properly.</p> <p>For Microsoft Cluster Services (MSCS) - When you need to stop I/O to a clone or a clone's source LUN, you must take the clone LUN offline in order to ensure MSCS stops generating I/O. You must do this in addition to using the <code>admsnap clone_deactivate</code> command. You can take the clone LUNs offline by doing the following:</p> <ol style="list-style-type: none"> 1. Locate the MSCS Disk Group resource that includes the clone LUNs. 2. Use the Microsoft Cluster Administrator to take the Disk Group offline by right-clicking the Disk Group and selecting Take Offline. 3. Perform the clone operation. <p>When you are ready to resume I/O to the disk, use the Microsoft Cluster Administrator to bring the Disk Group(s) back online by right-clicking the Disk Group and selecting Bring Online.</p>
AIX, HP-UX, and Solaris	<p>The <code>clone_deactivate</code> command is not supported; it is provided for script compatibility between Windows and UNIX only. If you issue the command on this server, <code>admsnap</code> will return an error.</p> <p>Unmount the file system by issuing the <code>umount</code> command. If you cannot unmount the file system, you can use the <code>sync</code> command to flush buffers. The <code>sync</code> command is not considered a substitute for unmounting the file system, but you can use it to reduce the number of times you have to issue the <code>fsck</code> command on your backup server's file system. Refer to your system's man pages for <code>sync</code> command usage.</p>
Linux	<p>The <code>clone_deactivate</code> command is not supported. You must either reboot the server or unload and load the HBA driver.</p>
NetWare	<p>The <code>clone_deactivate</code> command is not supported. Use the <code>dismount</code> command on the clone volume to dismount the file system.</p>

SYNTAX

```
admsnap clone_deactivate [-h] -o object_name [-t object_type]
```

OPTIONS

`-h`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

`-o object_name`

Specifies an object name. The object name must be a device name or drive letter. If the device name is provided as an object name, `admsnap` will find the associated device and drive letter. It will then delete the drive letter, dismount the device and take the device offline.

If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

`-t object_type`

Specifies the object type. Valid Windows object types are device and drive letter.

If you do not specify the `-t` option, `admsnap` will attempt to determine the correct object type and report a failure if it cannot.

EXAMPLE

Windows example:

```
admsnap clone_deactivate -o e:
```

admsnap flush

Flushes all data that has been written to the source LUN.

DESCRIPTION

The `admsnap flush` command clears buffers and flushes all data to a LUN.

Use it on the source LUN's server before fracturing a clone to ensure that all cached data has been written to the clone.

Depending on your operating system, the `admsnap flush` command will do the following:

Table 53 Function of `admsnap flush` command for an operating system

Server OS ^a	Function description
Windows	Writes the data from the buffers to the disk.
AIX, HP-UX, Linux, and Solaris	Writes the data from the buffers to the disk. We recommend that you unmount the file system by issuing the <code>umount</code> command to ensure that no additional source writes are in process. If you are unable to unmount the file system, you can issue the <code>admsnap flush</code> command ^b .
NetWare	The <code>flush</code> command is not supported. Dismount the file system by using the <code>dismount</code> command on the volume.

- a. With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements.
- b. Neither the `flush` command nor the `sync` command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

SYNTAX

```
admsnap flush [-h] -o object_name [-t object_type]
```

OPTIONS

`-h`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`-o object_name`

Specifies an object name. Valid Windows object names are drive letter and physical drive name. Valid Solaris object names are device names, file system names, or Veritas volume names. For HP-UX and Linux, a valid object name is either a device name, a file system name or Veritas volume name. For AIX, a valid object name is either a device name or a file system name.

If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

`-t` *object_type*

Specifies the object type. This switch is supported on Windows or UNIX operating systems only. Valid Windows object types are device and drive letter. Valid UNIX object types are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

If you do not specify the `-t` option, `admsnap` will try to determine the correct object type and will report a failure if it cannot.

EXAMPLE

Windows example:

```
admsnap flush -o E:
```

admsnap list

Performs a list operation and provides the World Wide Name (WWN) of all supported devices on the system.

PREREQUISITES

You must have administrative privileges on the production and secondary server.

DESCRIPTION

The `admsnap list` command will attempt to list devices and their World Wide Names (or WWN) on the operating system. You can use the WWN to distinguish one operating system device from another.

The `admsnap list` command is supported on Windows, Solaris, HP-UX, Linux, and AIX systems only.

SYNTAX

```
admsnap list [-h]
[-l lun_wwn]
[-d object_name]
[-a output_format]
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`[-l lun_wwn]`

Lists the device name for the WWN you specify. The WWN must be in a valid WWN format.

`[-d object_name]`

Specifies an object name that must be a device name, file system mount point, a Veritas volume name, or a drive letter.

Note

The `-a` option is ignored when used with the `-d` option and the object name argument is a Windows device name, for instance `\\.\PhysicalDrive1`.

[`-a output_format`]

Specifies an alternative output style on a Windows server. The default Windows output style is to list the drive letter and its WWN. The output style must be a valid style that is either a drive letter (default if `-a` is not specified at all) or a physical drive.

Note

The `-a` option is ignored when used with the `-d` option and the object name argument is a Windows device name, for instance `\\.\PhysicalDrive1`.

EXAMPLE

Windows example:

```
admsnap list -d E:
```

SnapView Snapshot CLI commands

This section explains the SnapView snapshot CLI commands that manage SnapView snapshots.

Setting up SnapView to use snapshots

This section describes how to set up SnapView to use snapshots.

Prerequisites for setting up snapshots

Before you can set up and use snapshots, the following must be true:

- Source LUNs must be bound. For a client or production server to access a source LUN, you must assign the source LUN to a storage group and connect the storage group to the production server. To do this, you must enable data access control on the system.
- For VMware ESX Servers, verify that the source LUN is presented to the virtual machine (guest operating system running on the virtual machine). For information on how to present a LUN to the virtual machine, refer to the VMware documentation that shipped with your ESX Server.
- For configuration restrictions when using VMFS volumes, go to the E-Lab Interoperability Navigator on EMC Online Support and under the **PDFs and Guides** tab, open the **VMware ESX server.pdf** file.
- For a secondary server to access the snapshot, you must assign the snapshot to a storage group. The storage group must be connected to the secondary server that will activate the snapshot. You must assign the snapshot to a storage group other than the storage group that holds the source LUN. EMC supports placing a snapshot in the same storage group as its source LUN only if you use Replication Manager or Replication Manager/SE to put the snapshot in the storage group. This software provides same host access to the snapshot and the source LUN. For information on using these software products, refer to the documentation for the product. If you have a VMware ESX Server, the snapshot and source LUNs must be accessed by different virtual machines, unless the virtual machine is running one of the software programs that supports same host access.

Note

VMware ESX Servers must activate the snapshot before adding it to a storage group.

- Reserved LUNs must be bound and added to the reserved LUN pool before starting a SnapView session. You must add a reserved LUN to the reserved LUN pool for each source LUN you will replicate. To configure the reserved LUN pool, refer to [VNX CLI for Block on page 23](#).
-

Note

Thin LUNs are not eligible for use as reserved LUNs.

Overview of setting up SnapView to use snapshots

The following is a checklist for setting up SnapView snapshots.

Note

You must complete the prerequisites for setting up snapshots, as listed in [Prerequisites for setting up snapshots on page 560](#), before you can perform any of the following procedures.

- Start a SnapView session, refer to [snapview -startsession on page 579](#).
-

Note

You can create a snapshot before starting a session but the snapshot has no use until you start a session on it. A secondary server can then activate the snapshot to the session.

- Create a snapshot, refer to [snapview -createsnapshot on page 566](#).
- If you do not have a VMware ESX Server - Add the snapshot to a storage group connected to the server that will access the snapshot.

If you have a VMware ESX Server - [Activate the snapshot on page 563](#). After you have activated the snapshot, do the following:

1. Add the snapshot to a storage group connected to the ESX Server that will access the snapshot.
2. Rescan the bus at the ESX Server level.
3. If a virtual machine is already running, power off the virtual machine and use the Service Console of the ESX Server to assign the snapshot to the virtual machine. If a virtual machine is not running, create a virtual machine on the ESX Server and assign the snapshot to the virtual machine.

If a virtual machine is not running, create a virtual machine on the ESX Server and assign the snapshot to the virtual machine.
4. Power on the virtual machine. If the snapshot is already presented to the virtual machine, rescan the bus at the virtual machine level. For virtual machines running Windows, you can use the `admsnap activate` command to rescan the bus.

Reserved LUN pool with SnapView

The reserved LUN pool consists of one or more private LUNs and works with SnapView sessions and snapshots. The reserved LUN pool stores the original source LUN data chunks that have been modified since the start of the session. For any one session, the contents of a reserved LUN(s) and any unchanged source LUN(s) blocks compose the snapshot.

Server writes made to an activated snapshot are also stored on a reserved LUN in the global reserved LUN pool. When you deactivate the snapshot, the reserved LUN space is freed and all snapshot writes are destroyed.

Before starting a session, the reserved LUN pool must contain at least one LUN for each source LUN that will be starting a session. You can add any LUNs that are available to the reserved LUN pool. Each system manages its own LUN pool and assigns a separate reserved LUN (or multiple LUNs) to each SnapView source LUN. Multiple sessions of a single source LUN will share the same reserved LUN or LUNs.

If the reserved LUN fills up and there are no available reserved LUNs in the reserved LUN pool, the software automatically terminates the session that is trying to allocate reserved LUN space, logs an error, releases the reserved LUN(s) used by this session, and returns them to the reserved LUN pool. The software also destroys all copy-on-first-write data stored in the reserved LUN pool for that session. At this point, the snapshot becomes inactive and any server that has mounted volumes on the snapshot will receive I/O errors and lose access.

If you have multiple sessions of a single source LUN and the reserved LUN fills up, when the production server modifies a chunk on the source LUN, resulting in a copy-on-first-write, every session that has the same chunk will be terminated if no additional LUNs are available in the reserved LUN pool. Other sessions that did not have this chunk will continue to run and use the reserved LUN space that the terminated sessions were using.

Note

SnapView, incremental SAN Copy, and MirrorView/Asynchronous share the LUN resources of the reserved LUN pool. For example, if you are running an incremental SAN Copy session on one LUN and a SnapView session on another LUN, the reserved LUN pool must contain at least two LUNs - one for each source LUN. If both sessions are running on the same source LUN, the sessions will share a reserved LUN.

Configuring the reserved LUN pool

You must configure the reserved LUN pool before you start a SnapView session. Use VNX CLI for Block to configure the reserved LUN pool (refer to [VNX CLI for Block on page 23](#)). The CLI Reference also explains how you can manually estimate a suitable LUN pool size.

Note

Thin LUNs are not eligible for use as reserved LUNs.

SnapView command functions

The navisecli or navicli snapview functions, listed alphabetically, are:

- `snapview -activatesnapshot`
- `snapview -chgrollback`
- `snapview -chgssnapshot`
- `snapview -createsnapshot`
- `snapview -deactivatesnapshot`
- `snapview -listrollback`
- `snapview -listsnapshots`
- `snapview -listsessions`
- `snapview -listsnapableluns`
- `snapview -lunpool`

- `snapview -rmsnapshot`
- `snapview -startrollback`
- `snapview -startsession`
- `snapview -stopsession`

In [naviseccli or navicli SnapView command functions on page 563](#), the SnapView snapshot CLI functions are listed alphabetically according to essential and optional functions.

Table 54 naviseccli or navicli SnapView command functions

Command	Description
Essential functions (alphabetically)	
<code>snapview -activatesnapshot</code> or <code>admsnap -activate</code>	Associates a session with a snapshot LUN, required for multiple sessions on a LUN(s). The <code>admsnap activate</code> command returns a drive letter or device name of the snapshot.
<code>snapview -createsnapshot</code>	Creates a snapshot LUN.
<code>snapview -startrollback</code>	Starts a rollback on the selected SnapView session.
<code>snapview -startsession</code>	Starts a single SnapView session.
<code>snapview -stopsession</code>	Stops a SnapView session.
Optional status functions (alphabetically)	
<code>snapview -listrollback</code>	Displays rollback information such as the rollback status and performance statistics.
<code>snapview -listsessions</code>	Displays information on all active SnapView sessions.
<code>snapview -listsnapableluns</code>	Displays the unique IDs of all potential snapshot source LUNs.
<code>snapview -listsnapshots</code>	Displays information on snapshot LUNs.
<code>snapview -lunpool</code>	Lists all reserved LUNs that are currently being used with SnapView.
Optional reconfiguration functions (alphabetically)	
<code>snapview -chgrollback</code>	Changes the rollback rate.
<code>snapview -chgsnapshot</code>	Renames a snapshot LUN.
<code>snapview -deactivatesnapshot</code>	Removes the association of a session with a snapshot LUN.
<code>snapview -rmsnapshot</code>	Removes (deletes) a snapshot LUN.

snapview -activatesnapshot

Activates a snapshot LUN with a session.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

A session must already exist (that is, you must start one using `snapview -startsession`) and a snapshot LUN must exist (`snapview -createsnapshot`) before this command will work.

If you have a VMware ESX Server, you must activate the snapshot before you add it to a storage group.

DESCRIPTION

The `naviseccli` or `navicli` `snapview` command with `-activatesnapshot` maps a snapshot LUN to a SnapView session. When you activate a snapshot to a session, the session becomes accessible on the secondary server's system.

Note

You must issue this command to the SP that owns the source LUN(s); otherwise, the command will fail.

Once a secondary server activates a snapshot to a session, this server can write to the activated snapshot. The software stores all writes made to the snapshot in the reserved LUN pool. If the secondary server deactivates the snapshot from the session, the software destroys all writes made to the session.

Note

The `snapview -activatesnapshot` command has a similar function as the `admsnap activate` command, but the `snapview -activatesnapshot` command does not return snapshot device information or allow you to activate all snapshots using one command.

SYNTAX

```

snapview -activatesnapshot sessionname -snapshotid uid
or
snapview -activatesnapshot sessionname -snapshotname name

```

OPTIONS

- `-activatesnapshot sessionname`
Specifies a session name (that you assigned when the session was started with `-startsession`).
- `-snapshotid uid`
Specifies the 32-byte unique ID of the snapshot (assigned by the system when the copy was created with `-createsnapshot`).
- `-snapshotname name`
Specifies the name of the snapshot (that you optionally assigned when the copy was created with `-createsnapshot`).

EXAMPLE

For SP `ss1_spa`, these commands start a session named `2pmsession` with snapshot LUN `lun20copy`, and then associate the session and snapshot. The copy then becomes accessible on the secondary server's system.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```

naviseccli -h ss1_spa snapview -startsession 2pmsession -snapshotname
lun20copy

```

```
naviseccli -h ss1_spa snapview -activatesnapshot 2pmsession
-snapshotname lun20copy
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -chgrollback

Changes the rollback rate.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-chgrollback` changes the rate at which the data is copied back to the source LUN(s).

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command from the SP that owns the source LUN(s) or the command will fail.

SYNTAX

```
snapview -chgrollback sessionname-rate [high|medium|low]
```

OPTIONS

`-chgrollback sessionname`

Specifies the name for the SnapView session that is rolling back.

`-rate high | medium | low`

Specifies the rate at which the software copies the data back to the source LUN(s). Options are `high`, `medium`, or `low`. The default is `medium`.

EXAMPLE

For SP `ss1_spa`, this command changes the rollback rate to `medium` for the session named `rollbackTest`. There is no output for this command. You can verify the new rollback rate by issuing the `-listrollback -rate` command. For additional information about the `-listrollback` command, refer to [snapview -listrollback on page 569](#).

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -chgrollback rollbackTest -rate medium
```

OUTPUT

If the specified session is not rolling back, the software returns an error message.

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -chgsnapshot

Renames a snapshot.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-chgsnapshot` changes the name of a snapshot.

Note

You can issue this command from either SP.

SYNTAX

```
snapview -chgsnapshot -snapshotid uid -snapshotnewname newname
or
snapview -chgsnapshot -snapshotname name -snapshotnewname
newname
```

OPTIONS

`-snapshotid uid`

Specifies the 32-byte LUN unique ID (WWN) of the snapshot, assigned when the snapshot was created. For any letters in the ID, you must specify the same case that the software assigned to the image when the snapshot was created (use the `snapview -listsnapshots` function).

`-snapshotname name`

Specifies the existing name for the snapshot.

`-snapshotnewname newname`

Specifies the new name for the snapshot.

EXAMPLE

For SP `ss1_spa`, this command renames the snapshot named `snap11am` to `snap3pm`.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -chgsnapshot -snapshotname snap11am
-snapshotnewname snap3pm
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -createsnapshot

Creates a snapshot.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-createsnapshot` creates a snapshot. A snapshot is a virtual LUN and when activated (`-activatesnapshot`), it allows a secondary server to view a SnapView session. An active snapshot is a composite of a source LUN and reserved LUN data that lasts until you destroy the snapshot. You can create a snapshot before or after you start a session; however, the snapshot has no use until a secondary server activates it to a session.

Note

When you create a snapshot, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

If the system loses power while the SnapView session is running, both the session and snapshot will survive the system power failure, since all sessions run in persistence mode.

You must issue this command to the SP that owns the source LUN(s); otherwise, the command will fail.

Note

Unless you have additional software that supports same host access, you must assign the snapshot to a storage group other than the storage group that holds the source LUN(s). You also must assign multiple snapshots, of the same source LUN(s), to different storage groups. For information on software that supports same host access, refer to the [Prerequisites for setting up snapshots on page 560](#).

SYNTAX

```
snapview -createsnapshot lun [-snapshotname name]
```

OPTIONS

`-createsnapshot lun`

Specifies the LUN ID of the source LUN(s).

`-snapshotname name`

Lets you assign a meaningful name to the snapshot; for example `2pmsnap`. You can use this name during the session to manage the snapshot. The name can include as many as 32 characters and spaces are allowed.

EXAMPLE

For SP `ss1_spa`, this `-createsnapshot` command creates a snapshot for the source LUN(s) that has the LUN ID 20. The snapshot name is `lun20snap`.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -createsnapshot 20 -snapshotname
lun20snap
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -deactivatesnapshot

Removes the association between a SnapView session and a snapshot LUN.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-deactivatesnapshot` unmaps a snapshot from a SnapView session and destroys any secondary server writes made to the snapshot. The snapshot becomes inaccessible on the secondary server's system, while the session continues.

Note

You must issue this command to the SP that owns the source LUN(s); otherwise, the command will fail.

SYNTAX

```
snapview -deactivatesnapshot -snapshotid uid [-o]
or
snapview -deactivatesnapshot -snapshotname name
[-o]
```

OPTIONS

- `-snapshotid uid`
Specifies the 32-byte unique ID of the snapshot (assigned when the copy was created with `-createsnapshot`).
- `-snapshotname name`
Specifies the name of the snapshot (optionally assigned when the copy was created with `-createsnapshot`).
- `-o`
Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command removes the existing association between the snapshot `lun20copy` and its session. The snapshot then becomes inaccessible on the secondary server's system.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -deactivatesnapshot -snapshotname
lun20copy
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -listrollback

Displays rollback information, such as the rollback status and performance statistics.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

Note

If a session spans across both SPs, the switches that provide statistical information will specify only the statistics for the SP in the `navicli` command.

The `naviseccli` or `navicli snapview` command with `-listrollback` and no other switches lists all rollback information. You can obtain more specific information with the switches described below.

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command to the SP that owns the source LUN(s) or the command will fail.

SYNTAX

```

snapview-listrollback [-name sessionname]
[-rate [-high|-medium|-low]]
[-complete]
[-incomplete]
[-reads]
[-writes]
[-spanreads]
[-spanwrites]
[-blocksread]
[-blockswritten]
[-rollbackpercent]

```

OPTIONS

`-name sessionname`

Specifies the name for the SnapView session that is rolling back. If you do not specify a session name, the software returns the names for all sessions that are rolling back on that SP.

`-rate high|medium|low`

Specifies the rate at which the software copies the data back to the source LUN(s). Values are high, medium, or low.

`-complete`

Displays the number of chunks that have been copied to the source LUN(s). This includes the number of chunks flushed and the number of chunks invalidated.

The number of chunks flushed is the number of reserved LUN chunks that have been copied from the session that is rolling back to the source LUN(s).

The number of chunks invalidated is the number of reserved LUN chunks that will not be copied to the source LUN(s). Server writes made to the source LUN(s) while

the rollback is in progress will take precedence over the data chunks being rolled back.

- incomplete
Displays the number of copy-on-first-write chunks that have not been copied back to the session's source LUN(s).
- reads
Displays the number of server read requests made to the source LUN(s) during the rollback operation.
- writes
Displays the number of server write requests made to the source LUN(s) during the rollback operation.
- spanreads
Displays the number of server read requests made to the source LUN(s) during the rollback operation, which spanned a partial chunk.
- spanwrites
Displays the number of server write requests made to the source LUN(s) during the rollback operation, which spanned a partial chunk.
- blocksread
Displays the number of 512-byte disk blocks that were read by the source LUN's server during the rollback operation.
- blockswritten
Displays the number of 512-byte disk blocks that were written by the source LUN's server during the rollback operation. Server writes made to the source LUN(s) while the rollback is in progress will overwrite the data being rolled back.
- rollbackpercent
Displays the percentage of data chunks that have been copied to the source LUN(s).

EXAMPLE

You can specify no switches or only those that you are interested in, as illustrated in the examples below.

naviseccli -h ssl_spa snapview -listrollback

Without any switches specified for SP `ssl_spa`, this command will display the statistics information for all sessions that are rolling back.

naviseccli -h ssl_spa snapview -listrollback -name

For SP `ssl_spa`, this command displays only the names of all sessions that are rolling back.

naviseccli -h ssl_spa snapview -listrollback -name sessionName

For SP `ssl_spa`, this command displays the statistics information for the specified session.

naviseccli -h ssl_spa snapview -listrollback -name -rate -complete

For SP `ssl_spa`, this command displays the rollback rate and the number of chunks completed for each session that is rolling back.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

OUTPUT

The following is a sample output that shows the devices controlled by one agent. Actual output varies depending on the switches you use with the `-listrollback` command.

```
Name of the session: rollbackTest
Session RollBack Rate:          low
Number of Chunks Complete: 1024
Chunks Flushed: 1024
Chunks Invalidated: 0
Number of Chunks Incomplete: 15293
Host Read Requests: 0
Host Write Requests: 0
Spanning Host Read Requests: 0
Spanning Host Write Requests: 0
Number of Blocks Read: 0
Number of Blocks Written: 0
Percent Rolled Back: 6.27566
```

If no session is rolling back, VNX CLI for Block displays the following message:

```
No snap session on this SP is in the process of rolling back.
```

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -listsessions

Lists session performance statistics.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-listsessions` and no other switches lists information about all SnapView sessions. With the `-name name` switch, it lists all information on the specified session.

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command to the SP that owns the source LUN(s) or the command will fail.

SYNTAX

```
snapview -listsessions [-name [sessionname]]
[-all]
[-cacheusage]
[-consistent]
[-cowwrites]
[-creationtime]
[-includepeer]
[-luns]
[-lunwrites]
[-persistence]
[-readfromcache]
[-readslun]
[-readssnapshot]
```

```
[-reserved]
[-sessionstate]
[-snapshotsname]
[-snapshotsuid]
[-tlunumber]
[-totalwrites]
[-writeslargerthansize]
```

OPTIONS

```
-name [sessionname]
    Specifies the name of the session to be described. If you omit -sessionname, the CLI
    displays all session names.
-all
    Lists all properties of a session without having to type in any other switch.
-cacheusage
    Displays the system's reserved LUN pool capacity in gigabytes and percentage of or
    total number of chunks used in this SnapView session.
-consistent
    Lists whether the session is running in consistent mode.
-cowwrites
    Displays the number of write requests to the reserved LUN pool that triggered a copy
    on first write.
-creationtime
    Displays the time and date that the session was created.
-includepeer
    Lists sessions on the peer SP and LUNs from sessions spanning both SPs.
-luns
    Lists the source LUNs used in the sessions.
-lunwrites
    Displays the total number of write requests from the source LUN(s).
-persistence
    Lists whether the session is running in persistence mode. VNX CLI for Block version
    6.24 or later enables persistence mode as the default.
-readfromcache
    Displays the number of chunks read from the reserved LUN pool.
-readslun
    Displays the total number of chunks read from the source LUN(s).
-readssnapshot
    Displays the number of chunks read from the snapshot.
-reserved
    Lists all reserved sessions. Reserved sessions are sessions used for another
    application such as SAN Copy and MirrorView/Asynchronous.
-sessionstate
    Displays the state of the session: Stopping, Normal, or Rolling Back.
-snapshotsname
    Lists the names of any active snapshots in this session.
-snapshotsuid
    Lists the 16-byte LUN unique IDs of the snapshots in this session.
-tlunumber
    Displays the LUN number of the source LUN(s).
-totalwrites
    Displays the sum of write requests to the source LUN(s) and to the reserved LUN pool
    in the SnapView session.
-writeslargerthansize
    Displays the number of write requests larger than the chunk size.
```

EXAMPLE

For SP `ss1_spa`, this command lists all sessions. A sample output follows.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -listsessions
```

OUTPUT

The following is a sample output that shows the devices controlled by one agent. Actual output varies depending on the switches you use with the `-listsessions` command.

```
Name of the session: rollbackTest
Number of read requests serviced by the snapview reserved LUN
pool: 65264
Total number of read requests on the snapview logical unit: 0
Number of reads from the Thin LUN: 4294902032
Number of writes requests in the session: 65264
Number of writes requests to reserved LUN pool that triggered a
COW: 0
Total number of writes requests on the snapview target logical
unit: 65264
Number of writes requests larger than the chunk size: 0
LUN Pool Capacity in GB: 6
Session Usage for This SP: 33.257
List of Thin LUNs: LUN 8 LUN 9
snap Logical Units UID
snap Logical Units Name
Session in persistence: YES
Session creation time: 03/21/03 06:56:28
Session state: Normal
```

If the version of VNX OE for Block software running on the SP does not support this command, a Command is not supported error message is printed to `stderr`.

snapview -listsnapshots

Displays information on snapshot copies.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-listsnapshots` and no other switches lists all information on all snapshot copies.

With `-snapshotname name` or `-snapshotid uid`, it lists all information about one snapshot. The command works only after a `snapview -createsnapshot` command has created one or more snapshot copies.

Note

You can issue this command from either SP.

SYNTAX

```
snapview -listsnapshots [-lun]
[-name]
[-reserved]
```

```
[-snapshotname [name]]  
[-snapshotid [uid]]  
[-state]
```

OPTIONS

```
-lun  
    Lists source LUNs.  
-name  
    Lists the SnapView session names.  
-reserved  
    Lists all reserved snapshots. Reserved snapshots are snapshots used for another  
    application such as SAN Copy and MirrorView/Asynchronous.  
-snapshotname [name]  
    Specifies the name of the snapshot assigned when the snapshot was created or the  
    session started. If you omit name, it displays all snapshot names.  
-snapshotid [uid]  
    Specifies the unique ID of the snapshot. If you omit uid, it displays the unique IDs  
    of all copies.  
-state  
    Lists the state of the copy: Active or Inactive.
```

EXAMPLE

For SP `ss1_spa1`, this command lists all information on snapshot copies. A sample output follows.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -listsnapshots
```

OUTPUT

The following is a sample output that shows the devices controlled by one agent. Actual output varies depending on the switches you use with the `-listsnapshots` command.

```
SnapView logical unit name: snapOnLun4  
SnapView logical unit ID: 60:06:01:67:51:09:00:00:CD:2E:  
76:86:FE:5D:D7:11  
Target Logical Unit: 4  
State: Active  
Session Name: sessionAtNoonOnLun4
```

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -listsnapableluns

Lists all source LUNs on which you can create a snapshot.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-listsnapableluns` lists all source LUNs that you can snap.

Note

You can issue this command from either SP.

SYNTAX

```
snapview -listsnapableluns
```

EXAMPLE

For SP `ss1_spa`, this command lists all LUNs ready for snapshot copying. A sample output follows.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -listsnapableluns
```

OUTPUT

The following is a sample output that shows the devices controlled by one agent. Actual output varies depending on the switches you use with the `-listsessions` command.

```
Name LUN 1
Name LUN 2
Name LUN 3
Name LUN 4
Name LUN 5
Name LUN 8
Name LUN 9
Name LUN 10
Name LUN 11
Name metaLun67
```

If the software running in the SP does not support the command, the CLI prints a `Command is not supported` error message to `stderr`.

snapview -lunpool

Lists all reserved LUNs that are currently being used with SnapView.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-lunpool` lists all information about the reserved LUNs that SnapView is currently using on the source LUNs owned by the directed SP.

SYNTAX

```
snapview -lunpool
```

EXAMPLE

This command lists all reserved LUNs on the system being used by SnapView. A sample output follows.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ssl_spa snapview -lunpool
```

OUTPUT

The following is a sample output that shows the devices controlled by one agent.

Target LUNs	AssociatedLUNPoolLUNs	LUNPoolLUNUsedPercent
4	20	3.562500
5	21	3.562500
1	16	0.347900
0	17	0.347900

If the software running in the SP does not support the command, the CLI prints a Command is not supported error message to stderr.

snapview -rmsnapshot

Deletes a snapshot.

PREREQUISITES

For naviseccli, you must have a user account on the system on which you want to execute the command. For navicli, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The naviseccli or navicli snapview command with `-rmsnapshot` deletes a snapshot created with the `snapview -create` function.

You can specify either the snapshot unique ID or name (if you assigned a name when you created the copy).

Note

You must issue this command to the SP that owns the source LUN(s); otherwise, the command will fail.

SYNTAX

```
snapview -rmsnapshot -snapshotid uid [-o]
or
snapview -rmsnapshot -snapshotname name
[-o]
```

OPTIONS

- `-snapshotid uid`
Specifies the 32-byte unique ID of the snapshot (assigned when the copy was created with `-createsnapshot`).
- `-snapshotname name`
Specifies the name of the snapshot (optionally assigned when the copy was created with `-createsnapshot`).
- `-o`
Removes the copy without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, this command starts removing the snapshot with the name `2pmsnap`. The software asks for confirmation.

Note

To use this command with `navicli`, replace `navisecli` with `navicli`.

```
navisecli -h ssl_spa snapview -rmsnapshot -snapshotname 2pmsnap
```

```
Do you really want to remove the snapshot (y/n) [N]? y
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

snapview -startrollback

Starts a rollback on the selected SnapView session.

PREREQUISITES

For `navisecli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

Before starting a rollback on a SnapView session, verify the following:

- The session is running in persistence mode. VNX CLI for Block version 6.24 or later enables persistence mode as the default.
- The session is not in the process of being rolled back.
- The source LUN(s) is not participating in another rollback.
- If an activated snapshot is mapped to this session, unmount the snapshot.

DESCRIPTION

Note

Once you start a rollback, you cannot stop it or the session that is rolling back.

The `navisecli` or `navicli snapview` command with `-startrollback` starts a rollback on the selected SnapView session. Rollback lets you restore the point-in-time copy of a session to the source LUN(s).

CAUTION

Windows operating systems only - To prevent data corruption during the rollback operation, you should disable the indexing service and recycle bin on the source LUN(s) of the session you will roll back.

When you confirm the start of a rollback operation, the source LUN can instantly access the session's point-in-time data, while background copying continues (until all the session's point-in-time blocks are copied to the source LUN). You cannot start a new session on any source LUN that is participating in the rollback operation until all its background data copying has completed. The rollback operation itself does not count against the eight-session limit per source LUN.

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command to the SP that owns the source LUN(s) or the command will fail.

Rolling back with snapshots

Each source LUN belonging to a session that is rolling back can have an activated or deactivated snapshot. If the session has an activated snapshot, the software copies any server writes made to this snapshot to the source LUN. However, you must unmount the snapshot before you start the rollback operation. When the rollback operation completes, including background copying, you can remount the snapshot.

If you deactivate the snapshot prior to starting the rollback operation, any server writes made to the snapshot are lost.

Allocating reserved LUN pool space

The rollback operation itself does not use reserved LUN pool space. However, the rollback may cause copy-on-first-write operations for sessions you started after the session that is rolling back. This in turn will consume more reserved LUN pool space. These copy-on-first-write operations occur because the original point-in-time data of those sessions is being modified by the rollback operation. Additionally, since you can resume I/O to the source LUN(s) or start another session while background copying continues, you may need additional reserved LUN space.

If you run out of reserved LUNs in the SP's LUN pool while the rollback operation is in progress (during the background copying), the software will terminate the session that is rolling back after it completes all background copying. The software will also terminate any other session that tries to allocate additional reserved LUN pool space.

You cannot roll back a SnapView session if the session's source LUN(s) is:

- an unfractured clone. You must fracture the clone before rolling back the session.
- a fractured clone that is unavailable for I/O, for instance, if the clone was fractured while in an out-of-sync state.
- a clone group's source LUN that has any unfractured clone LUNs or clones that are in the process of reverse synchronizing. All clone LUNs for this source LUN must be fractured in a synchronized or consistent state before you can roll back a session to its source LUN.
- a primary image that has one or more unfractured secondary images. You must fracture the secondary image(s) before rolling back the session.
- a secondary mirror image. You must promote the secondary image and then fracture it before rolling back the session.

SYNTAX

```
snapview -startrollback sessionname [-rate high|medium|low]
[-o]
```

OPTIONS

`-startrollback sessionname`

Specifies the name for the SnapView session you want to roll back.

`-rate high | medium | low`

Specifies the rate at which the software copies the data back to the source LUN(s). Options are `high`, `medium`, or `low`. The default is `medium`. You can change this rate while the rollback operation is in progress (see [-chgrollback on page 565](#)). If you do not specify the `-rate` option, the software uses the default value of `medium`.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command starts rolling back the session named `rollbackTest`. The specified rollback rate is `low`. The software asks for confirmation.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ssl_spa snapview -startrollback rollbackTest -rate low
```

This session has at least one active snapshot associated with it. Continuing this operation will restore not just the original data but also any modifications made to the active snapshot LUNs. If only the original data is desired for the RollBack, then quit this operation and deactivate the snapshot first.

Make sure that you have quiesced all host I/O to the Session's source LUNs and to any of the associated active Snapshot LUNs. This command will flush the Session's data back to the Source LUNs asynchronously in the background.

NOTE: If you want a recovery session for this rollback, you must create a persistent session NOW before beginning the rollback. Otherwise you will not be able to recover the current data once rollback begins.

Do you want to continue? (y/n)?

OUTPUT

If the version of software running on the SP does not support this command, a Command is not supported error message is printed to `stderr`.

snapview -startsession

Starts a SnapView session in persistence mode, and optionally in consistent mode.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a `username/server` entry in the Unisphere SP Agent configuration file.

Before starting a SnapView session, verify the following:

- You configured the reserved LUN pool by using the reserved `-lunpool -addlun` function.
- The source LUNs are currently not being used to start another consistent session. Once the consistent session starts on all the source LUNs, you can start another consistent session using the same source LUNs.
- If MirrorView/Asynchronous is installed, the source LUNs that you will use to start the consistent session cannot be in the process of starting an update to a group. If they are, you may receive an error when you attempt to start the consistent session. Once the mirror or group is updating or has been updated and before the next update occurs, you can start the consistent session using the same source LUNs.

DESCRIPTION

Note

With VNX CLI for Block version 6.24 or later, all SnapView sessions run in persistence mode, so the `-persistence` switch is no longer required.

The `naviseccli` or `navicli snapview` command with `-startsession` starts a SnapView session on the specified source LUN(s). A SnapView session is a point-in-time copy of a source LUN. The session keeps track of how the source LUN looks at a particular

point in time. You start a session using the snapshot name, snapshot unique ID (snapshot WWN), or the source LUN's number.

With VNX CLI for Block version 6.24 or later, all SnapView sessions run in persistence mode, which means that the session will survive SP reboot or failure, software installation, system reboot or power failure, and server I/O trespassing to the peer SP. You can also run a session in consistent mode using the `-consistent` option, which will preserve the point-in-time restartable copy across a set of source LUNs.

Note

When you start a session, SnapView automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

After you start a SnapView session and as the production server writes to the source LUN(s), the software stores a copy of the original data in the reserved LUN pool in chunks. This copy is referred to as copy-on-first-write and occurs only once, which is when the server first modifies a data chunk on the source LUN(s).

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command to the SP that owns the source LUN(s) or the command will fail.

SYNTAX

```
snapview -startsession sessionname -snapshotid uids |
-snapshotname names | -lun luns [-consistent] -persistence
```

OPTIONS

`-startsession sessionname`

Specifies the name of the session to be started. A session name can have as many as 32 characters. A name is required. If you use spaces in the name, enclose the entire name string in quotation marks in this command and in other CLI commands that access the session.

`-snapshotid uid`

Specifies one or more snapshot 32-byte unique IDs (WWNs). The snapshot software assigns a unique ID to the snapshot when you create it with `-createsnapshot`. For multiple IDs, separate IDs with blanks. For any letters in the ID, you must specify the same case as the software assigned to the image when the snapshot was created (use the `snapview -listsnapshots` function).

`-snapshotname names`

Specifies one or more snapshot names that you assigned when you created the snapshot with `-createsnapshot`. For multiple names, separate the names with commas.

`-lun luns`

Specifies one or more source LUN numbers. For multiple LUNs, separate the numbers with blanks.

`-consistent`

Specifies a consistent session. The session will preserve the point-in-time restartable copy across a set of source LUNs. The SnapView driver will delay any I/O requests to the set of source LUNs until the session has started on all LUNs (thus preserving the point-in-time restartable copy on the entire set of LUNs).

Note

A restartable copy is a data state having dependent write consistency and where all internal database/application control information is consistent with a database management system/application image.

If the session spans both SPs you must specify only one SP and the software will start the session on both SPs. Once you start a consistent session, you cannot add another source LUN to the session. In the event of a failure, the software will not start the session on any source LUN and will display an error message.

You would use consistent mode if any of the following is true:

- You want a consistent replica across a set of LUNs.
- You do not want to stop I/O to the source LUN(s) before starting a session.
- You want to prevent other LUNs from being added to the session. If desired, you can start a consistent session on a single source LUN to prevent other LUNs from being added to the session.
- The set of source LUNs that you will use to start the session spans both SPs; however, the source LUNs cannot span systems.

The consistent feature is available on a per-session basis (not per snapshot or source LUN) and counts as one of the eight sessions per source LUN limit.

Note

When a consistent session is in the process of starting on all LUNs, you cannot stop the session (the software may stop the session immediately after the start completes if I/O is queued and there is not enough space in the reserved LUN pool), perform a rollback operation, or activate a snapshot to the session.

`-persistence`

Note

With VNX CLI for Block version 6.24 or later, all SnapView sessions run in persistence mode, so the `-persistence` switch is no longer required.

Specifies a persistent session. The session will survive SP reboot or failure, software installation, system reboot or power failure, and server I/O trespassing to the peer SP.

The persistence feature is available on a per-session basis (not per snapshot or source LUN). In the event of a failure, reserved LUNs along with the source LUNs will trespass to the other SP. Depending on your failover software, once the failed SP is running, you may need to issue a restore command in order to restore the proper source LUNs and reserved LUNs back to their original SP. For the appropriate restore command, refer to the documentation that shipped with your failover software.

Note

For information on supported failover software for the system you are managing, refer to the *EMC SnapView and Admsnap Release Notes*.

EXAMPLE

In the first example, the `-startsession` command starts a persistent session named `2pmsession`. In the second example, the `-startsession` command starts a persistent and consistent session named `2pmsession` on LUNs one, two, and three. If

the session spans both SPs you need to specify only one SP and the software will start the session on both SPs.

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ssl_spa snapview -startsession 2pmsession
```

```
naviseccli -h ssl_spa snapview -startsession 2pmsession -persistence
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`.

POSTREQUISITES

After you start a SnapView session, a secondary server can then activate (map) a snapshot to the SnapView session. The snapshot views the original source LUN data chunks that have been modified since you started the session from the reserved LUN pool and unmodified data chunks from the source LUN(s).

snapview -stopsession

Stops a SnapView session.

PREREQUISITES

For `naviseccli`, you must have a user account on the system on which you want to execute the command. For `navicli`, you must have a username/server entry in the Unisphere SP Agent configuration file.

DESCRIPTION

The `naviseccli` or `navicli snapview` command with `-stopsession` stops the specified session and frees the reserved LUN pool space used by the session.

If the snapshots participating in the session belong to one or more storage groups and you stop the session, the servers connected to the storage groups will no longer have access to the snapshots in those storage groups. EMC recommends that you do one of the following:

- Flush the I/O on the server viewing the active snapshot before stopping the session.
- Deactivate the snapshot before stopping the session.

Stopping a session with an active snapshot makes the snapshot appear inactive.

Note

If a session spans across both SPs or if the session is a consistent session, you can issue this command from either SP; otherwise, you must issue this command to the SP that owns the source LUN(s) or the command will fail. If you started a session on multiple source LUNs, you can select any of the source LUNs to stop the session.

SYNTAX

```
snapview -stopsession sessionname [-o]
```

OPTIONS

```
-stopsession sessionname
```

Specifies the name of the session to be stopped. Someone must have assigned the name when the snapshot was created or when the session was started.

```
-o
```

Stops the session without asking for confirmation.

EXAMPLE

For SP `ss1_spa`, this command stops the session named `2pmsession`. The software asks for confirmation:

Note

To use this command with `navicli`, replace `naviseccli` with `navicli`.

```
naviseccli -h ss1_spa snapview -stopsession 2pmsession
```

```
Stopping of the session will stop data SnapView.
Do you really want to stop the session (y/n) [n]? y
```

OUTPUT

If the version of software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. Other errors:

```
INVALID_SESSION Session does not exist.
```

POSTREQUISITES

For a step-by-step overview and examples on setting up and using snapshots with `admsnap` and the CLI, refer to SnapView examples in *EMC VNX for Block Concepts*.

Admsnap commands for snapshots

This section explains how to use the `admsnap` utility software with SnapView snapshots.

About admsnap software and snapshots

The `admsnap` utility is an executable program that you can run interactively or with a script to manage snapshots. The `admsnap` utility resides on the servers connected to the system with the SnapView driver. You must execute the following `admsnap` commands from one of these attached servers:

- `admsnap activate` – activates a snapshot to a session.
- `admsnap deactivate` – deactivates a snapshot from a session.
- `admsnap flush` – flushes cached data to the LUN.
- `admsnap list` – lists all current snapshots.
- `admsnap start` – starts a session.
- `admsnap stop` – stops a session.

Note

The `admsnap` software supports native device names and PowerPath pseudo-device names. If you use a PowerPath pseudo-device name, you can add this device name to the `etc/fstab` file (for Linux) or to the `etc/vfstab` file (for Solaris). `Admsnap` supports both device naming scheme and enclosure naming scheme for Veritas.

For a step-by-step overview and examples of setting up and using snapshots with `admsnap` and VNX CLI for Block, refer to *SnapView examples* in *EMC VNX for Block Concepts*.

Reserved LUN pool (snapshot cache)

With SnapView version 02.03.xxx (or higher), the snapshot cache is referred to as the reserved LUN pool. The reserved LUN pool and the snapshot cache work with SnapView in

the same way. However, unlike the snapshot cache, which was used solely for SnapView, the reserved LUN pool shares its LUN resources with other applications such as SAN Copy and MirrorView/Asynchronous. VNX CLI for Block snapshot cache commands are still supported; however, EMC recommends that you use the new reserved LUN pool commands. For information on the new reserved LUN pool commands, refer to [VNX CLI for Block on page 23](#).

For reserved LUN pool information that is specific to SnapView, refer to [SnapView command functions on page 562](#).

Requirements for snapshots

Before you can use `admsnap` for snapshots, the following must be true:

General requirements

- A connected server must be running a suitable web browser to run VNX CLI for Block (and optionally the Unisphere). CLI is needed to configure and maintain the reserved LUN pool.
- The system must have the SnapView software enabled.
- The production and secondary servers must have both the Unisphere Host Agent and `admsnap` software installed. To install the host agent and `admsnap`, refer to the EMC Block server support products guide for your operating system.
- A reserved LUN must be available for each source LUN(s) for which you will start a session; otherwise an `admsnap start` command will fail. You can bind a reserved LUN and add it to an SP's reserved LUN pool using VNX CLI for Block or Unisphere.

Platform/server-specific requirements

- For HP-UX, `admsnap` requires a volume group entry for each path to the system. The entry or entries can be contained in a single volume group or multiple volume groups. On the secondary server, the volume group entry or entries must specify a snapshot LUN.
- For UNIX, if a session is unexpectedly stopped (due to SP panic or path failure), the secondary server must unmount the inactive snapshot or reboot the secondary server in order for this server to disregard the session's point-in-time data.
- For Windows, you may have to reboot the secondary server the first time you add the snapshot to the server's storage group.
- If you have a VMware ESX Server, you must activate the snapshot before you add the snapshot to a storage group. Refer to the [Overview of setting up SnapView to use snapshots on page 561](#) for additional information.

`admsnap activate`

Allows an attached server to access a snapshot.

PREREQUISITES

You must have administrative privileges on the secondary server.

DESCRIPTION

From the secondary server, the `activate` command scans system buses for system devices and determines if any device is part of a SnapView session.

When a secondary server activates (maps) a snapshot to a session, this server can write to the snapshot. The software stores all writes made to the snapshot in the reserved LUN pool. If the secondary server deactivates the snapshot from the session, the software destroys all writes.

Note

If you failed to flush the file system buffers before starting the session, the snapshot may not be usable.

You can activate only one session to a snapshot. To list current sessions use the CLI command `naviseccli snapview -listsessions`.

Note

The `snapview -activatesnapshot` command serves a similar function to the `admsnap activate` command, but does not return snapshot device information as `admsnap` does or allow you to activate all snapshots using one command. You must issue the `admsnap activate` command from the secondary server (the server that will access the snapshot).

Depending on your operating system, the `admsnap activate` command will do the following:

Table 55 Operating system and `admsnap activate` command

Server OS	Function description
Windows	<p>Returns the snapshot drive letter. You can use this drive immediately.</p> <p>If the activated snapshot has multiple volumes or partitions, <code>admsnap</code> will assign a drive letter to each volume or partition. If the volume is a dynamic drive, <code>admsnap</code> will activate all the physical drives but will not assign a drive letter. Additional steps will be required to bring the volume online and to assign the drive letter. If the given session is on multiple physical disks and each physical disk has its own volume, then each volume will be assigned a drive letter. If no drive letters are available, then <code>admsnap</code> will activate only one physical disk and an error appears through the command line.</p> <p>For Windows Remote Desktop and Terminal Services - If a terminal server is running on the secondary server, you must remove and disable it before using the <code>activate</code> command. If you have the terminal server enabled, it prevents the <code>activate</code> command from working properly.</p>
AIX, HP-UX, Linux, and Solaris	<p>Returns the snapshot device name. You will need to run <code>fsck</code> on this device only if it contains a file system and you did not unmount the source LUN(s) when you started the session. Then, if the source LUN(s) contains a file system, mount the file system on the secondary server using the snapshot device name to make the file system available for use.</p>

SYNTAX

```
admsnap activate [-h] -s session_name [-o object_name] [-t object_type] | [-d driveletter] [-r 0 | -r 1] [-c rescan_count]
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message will be displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`-s sessionname`

Specifies the session name that you assigned when you started the session. This is a string of up to 256 characters, and can be composed of letters, numbers and underscores.

The session name is case sensitive and must be unique on the system. If the session name contains spaces, you must enter the name in double quotation marks (for example, "June 21st Session").

`[-o object_name]`

Specifies an object name that you want to activate. You should use this option only when `admsnap` can determine the snapshot device list for the specified object name.

In UNIX, the objects that are generally available for this option are the raw UNIX device names and the EMC PowerPath device names. In Windows, the objects that are generally available for this option are the raw Windows device names.

If you omit this switch, the software searches all devices for the appropriate device, which takes time. You cannot use this option with the `-d` option. If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

The table below lists the supported object names for your operating system.

Table 56 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile views) • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Solaris	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Linux	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c • Linux MPIO's mpath device
AIX ^a	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device

- a. If the object name is on a Logical Volume Manager, then `admsnap` will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.

Table 56 Operating system object names (continued)

- b. If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- c. If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

[-t *object_type*]

Specifies the object type. Valid Windows object types are device, file system, volume, and drive letter. Valid UNIX object types are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

If you do not specify an object type, admsnap will attempt to determine the correct object type among the valid object types and report a failure if it cannot.

You should use the `-t` option when the argument for the `-o` option is not in a standard and recognizable form for your operating system's object or device. This means if you specify a valid argument for the `-o` option, you do not need the `-t` option for any operating system drive letter, file system, Veritas volume, or PowerPath pseudo-device names. For example, on Solaris systems, admsnap determines the `-o` argument of `/dev/rdisk/cXdXtXsX` to be a Solaris native device name, so you do not need to specify the `-t device`.

You cannot use this option in conjunction with the `-d` option.

[-d *driveletter*]

Specifies the drive letter you want to assign to the snapshot device, and if available, admsnap will assign the specified drive letter to the snapshot device. If the operation resulted in more than one device being activated, then all devices are activated but only the first device is assigned the drive letter.

Note

The `admsnap activate` command does not allow the use of drive letters A:, B:, C:, or D: with the `-d` option.

This option is valid on Windows systems only. All other operating systems ignore this option. You cannot use this option in conjunction with the `-o` or `-t` options.

[-r 0 | -r 1]

Specifies whether to perform a rescan of the operating system devices before performing the `activate` operation. The default rescanning behavior (when the `-r` option is not present on the command line) depends on whether you specify an object name with the `-o` option. If you do specify the object name, the default is not to perform a rescan. If you do not specify the object name, the default is to perform a rescan, which may take some time because it will rescan every device on the operating system.

When you use the `-r [0|1]` option, it overrides any default behavior.

This option is valid on Windows, Solaris, HP-UX, and Linux systems only. All other operating systems ignore this option.

[-c *rescan_count*]

Specifies the number of rescans for devices on the operating system. A rescan count occurs after a successful `activate` operation to bring the snapshot devices online. If you specify a count of 0, admsnap will not perform a rescan; however, the device could activate successfully on the system but will not be online with the server. If you omit the `-c` option, then a default value of 1 is set.

This option is valid on Windows only. All other operating systems ignore this option.

EXAMPLE

Windows example:

```
admsnap activate -s session1 -d e:
```

admsnap deactivate

Allows the secondary server to remove access to a snapshot and destroys any writes the secondary server made to the snapshot.

PREREQUISITES

You must have administrative privileges on the secondary server.

DESCRIPTION

The `admsnap deactivate` command unmaps a snapshot LUN from a SnapView session by removing access to the snapshot device from the secondary (backup) server. The deactivated snapshot goes offline and thereby becomes inaccessible to the secondary server.

Use this command on the secondary (backup) server after the session has been activated and you no longer require the snapshot device.

Note

The `snapview -deactivatesnapshot` command has a similar function as the `admsnap deactivate` command, but does not remove the snapshot device information from the server. You must issue the `admsnap deactivate` command from the secondary server (the server that activated it).

When the secondary server deactivates the snapshot from the session, `admsnap` destroys all writes made to the snapshot.

Depending on your operating system, the `admsnap deactivate` command will do one of the following:

Table 57 Operating system and `admsnap deactivate` command

Server OS	Function description
Windows	<p>Flushes all buffers and removes the drive letter assigned with <code>admsnap activate</code>. The snapshot goes inactive and continued I/O is not possible. This command lets you avoid rebooting the secondary server between activates.</p> <p>If the deactivated snapshot has multiple volumes or partitions, <code>admsnap</code> will remove a drive letter from each volume or partition. If the volume is a dynamic drive, <code>admsnap</code> will deactivate all the physical drives but will not remove the drive letter. Additional steps will be required to bring the volume offline and to remove the drive letter. If the given session is on multiple physical disks and each physical disk has its own volume, then the drive letter on each volume will be removed.</p> <p>For Windows Remote Desktop and Terminal Services - If a terminal server is running on the secondary server, you must remove and disable it before using the <code>deactivate</code> command. If you have the terminal server enabled, it prevents the <code>deactivate</code> command from working properly.</p>
AIX, HP-UX, Linux, and Solaris	Removes access to the snapshot but does not flush buffers.

Table 57 Operating system and admsnap deactivate command (continued)

Server OS	Function description
	Before you deactivate the snapshot, you must unmount the file system by issuing the <code>umount</code> command to flush all buffers. If you are unable to unmount the file system, you can issue the <code>admsnap flush</code> command. ^a

- a. The `flush` command is not considered a substitute for unmounting the file system; it only complements unmounting the file system.

SYNTAX

```
admsnap deactivate [-h] -s session_name [-o object_name [-t
object_type]] [-f]
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. Admsnap will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`-s sessionname`

Specifies the session name that you assigned when you started the session. This is a string of up to 256 characters, and can be composed of letters, numbers and underscores.

The session name is case sensitive and must be unique on the system. If the session name contains spaces, you must enter the name in double quotation marks (for example, "June 21st Session").

The name you specify must be a valid session name for an active SnapView session started from a production server and activated on the secondary (or backup) server that is performing the deactivate operation. All snapshots with the given session name will be deactivated.

On Windows systems, if the activated LUN with the given session name has multiple volumes and the system assigned drive letters, then admsnap will delete each drive letter, dismount the volumes, and then set the volumes offline before deactivating the LUN. If multiple LUNs have the same given session, then admsnap deletes the drive letters for all volumes and sets the volumes offline before deactivating the LUNs.

On UNIX servers, admsnap only deactivates the devices and no other operations are performed before deactivating the device.

`[-o object_name]`

Specifies an object name. When you execute the `admsnap deactivate` command with a device name as an object name, admsnap attempts to deactivate the snapshot on that device with the given session name. If the session name does not exist on the device, admsnap will report an error.

If you specify an object name that contains multiple devices, admsnap will deactivate the session on all the LUNs that make up those devices.

If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

The table below lists the supported object names for your operating system.

Table 58 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile views) • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Solaris	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Linux	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c • Linux MPIO's mpath device
AIX ^a	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device

- If the object name is on a Logical Volume Manager, then admsnap will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.
- If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

[*-t object_type*]

Specifies the object type. This option is valid only when the *-o* option is given. Valid Windows object types are device, file system, volume, and drive letter. Valid UNIX object types are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

If you do not specify an object type, admsnap will attempt to determine the correct object type among the valid object types and report a failure if it cannot.

This switch is supported on Windows or Solaris operating systems only.

[*-f*]

You can use this option to force a deactivate operation on a physical device name that does not have a drive letter assigned.

EXAMPLE

Windows example: Example of a fully qualified pathname to a Veritas volume named scratch:

```
admsnap deactivate -s session1 -o e:
```

Example of a fully qualified pathname to a raw Veritas device name:

```
admsnap deactivate -o /dev/vx/dsk/scratchdg/scratch
```

```
admsnap deactivate -o /dev/vx/rdmp/clt0d0s2
```

admsnap flush

Flushes all data that has been written to the source LUN.

PREREQUISITES

You must have administrative privileges on the production server.

DESCRIPTION

The `admsnap flush` command flushes all data.

Use it on the production server before starting a SnapView session to ensure that all cached data has been written to disk. Also, use the `flush` command on the secondary server before issuing a `deactivate` command to ensure that all data has been written.

Depending on your operating system, the `admsnap flush` command will do the following:

Table 59 Function of `admsnap flush` command for an operating system

Server OS ^a	Function description
Windows	Writes the data from the buffers to the disk.
AIX, HP-UX, Linux, and Solaris	Writes the data from the buffers to the disk. We recommend that you unmount the file system by issuing the <code>umount</code> command to ensure that no additional source writes are in process. If you are unable to unmount the file system, you can issue the <code>admsnap flush</code> command ^b .

- a. With some operating systems, you may need to shut down the application to flush the data. Specific operating systems have different requirements.
- b. Neither the `flush` command nor the `sync` command is a substitute for unmounting the file system. Both commands only complement unmounting the file system.

SYNTAX

```
admsnap flush [-h] -o object_name [-t object_type]
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`-o object_name`

Specifies an object name. If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

The table below lists the supported object names for your operating system.

Table 60 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile views) • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Solaris	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Linux	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c • Linux MPIO's mpath device
AIX ^a	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device

- If the object name is on a Logical Volume Manager, then admsnap will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.
- If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

[`-t object_type`]

Specifies the object type. This switch is supported on Windows or Solaris operating systems only. Valid Windows object types are device and drive letter. Valid Solaris object types are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

If you do not provide an object type, admsnap will attempt to determine the correct object type from the supported types and report a failure if it cannot.

You should use the `-t` option when the argument for the `-o` option is not in a standard and recognizable form for your operating system's object or device. This means if you specify a valid argument for the `-o` option, you do not need to specify the `-t` option for any operating system drive letter, file system, Veritas volume, or native or PowerPath pseudo-device names. For example, on Solaris systems, admsnap determines the `-o` argument of `/dev/rdisk/cXdXtXsX` to be a Solaris native device name, so you do not need to specify the `-t device`.

EXAMPLE

Windows example:


```
admsnap flush -o e:
```

admsnap list

Performs a list operation and provides the World Wide Name (WWN) of all supported devices on the system.

PREREQUISITES

You must have administrative privileges on the production and secondary server.

DESCRIPTION

The `admsnap list` command will attempt to list devices and their World Wide Names (or WWN) on the operating system. You can use the WWN to distinguish one operating system device from another.

The WWN is a worldwide unique string that the system provides for its devices. Use the WWN to uniquely identify those devices on the server. Use this operation to match the server's devices with the devices the server sees on the system.

This command is supported on Windows, Solaris, HP-UX, Linux, and AIX operating systems only.

SYNTAX

```
admsnap list [-h] [-l lun_wwn] [-d object_name] [-a output_format]
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. Admsnap will not perform the command; it lists the help output only.

This option is valid for Windows, and Solaris systems only.

`[-l lun_wwn]`

Lists the WWN of all supported devices on the system. The WWN must use a valid WWN format.

`[-d object_name]`

Specifies an object name. On Windows systems, the `-a` option is ignored when used with the `-d` option and the object name argument is a Windows physical device name, for example: `\\.\PhysicalDrive1`.

The table below lists the supported object names for your operating system.

Table 61 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile) • File system name^b • Veritas volume name (disk-based and enclosure-based)^c • Powerpath device
AIX ^a	<ul style="list-style-type: none"> • Device name

Table 61 Operating system object names (continued)

Operating system	Supported object name(s)
	<ul style="list-style-type: none"> • File system name^b • Powerpath device
Linux	<ul style="list-style-type: none"> • Device name • File system name • Powerpath device • Linux MPIO's mpath device • Veritas volume (disk-based and enclosure-based)
Solaris	<ul style="list-style-type: none"> • Device name • File system name • Powerpath device • Veritas volume (disk-based and enclosure-based)

- If the object name is on a Logical Volume Manager, then admsnap will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.
- If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

`[-a output_format]`

Specifies an alternative output style on a Windows server. The default Windows output style is to list the drive letter or physical drive and its WWN. Valid output formats are drive letter (default if `-a` is not specified at all) or physical drive. If the output format contains spaces, enter it without spaces.

Admsnap will ignore the `-a` option when you specify the `-d` option and the object name argument is a Windows device name, for example `\\.\PhysicalDrive1`.

EXAMPLE

Windows example:

```
admsnap list -d \\.\PhysicalDrive1
```

admsnap start

Starts a point-in-time copy of a source LUN.

PREREQUISITES

You must have administrative privileges on the production server.

DESCRIPTION

The `admsnap start` operation starts a SnapView persistent session with the specified session name. A persistent session survives SP failures, software installation upgrades, and event trespasses.

Note

Admsnap version 2.24 or later enables all SnapView sessions to run in persistence mode as the default.

From the production server, the `start` command specifies systems that have a LUN or file system you want to copy, then it begins the point-in-time copy of the specified source.

Note

Before you start a SnapView session, `admsnap` requires that you create a snapshot using the `-createsnapshot` command.

The `admsnap start` command is valid on Windows, Solaris, HP-UX, Linux, NetWare, and AIX systems.

SYNTAX

```
admsnap start [-h] -s session-name -o object-name [-t
object_type] -p -c
```

OPTIONS

`[-h]`

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

`-s session_name`

Specifies the session name that you assigned when you started the session. This is a string of up to 256 characters, and can be composed of letters, numbers and underscores.

The session name is case sensitive and must be unique on the system. If the session name contains spaces, you must enter the name in double quotation marks (for example, "June 21st Session").

Note

If you issue multiple start operations from the production server (this includes other production servers) using the same session name, the session will consist of all the LUNs that make up those devices. This operation can be called multiple times (even from different servers or EMC Unisphere) with the same session name and the result is that the new objects are added into the existing session.

`-o object_name`

Specifies an object name. You must specify an object name to identify the session. If you specify a device name or a drive letter, the LUN(s) you specify may have multiple partitions; if so, the snapshot will consist of all partitions on that LUN(s).

If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

Note

Admsnap version 2.1 with a Windows, or UNIX server - If you specify an object name that spans multiple devices with the `admsnap start` command, and one of those devices fails to start, admsnap will stop the session on all the devices. For example, you start a session on a Solaris file system and that file system spans across two LUNs on the system. The first device successfully starts the session but the second does not. The admsnap software will stop the first session started on the device, and then you must restart the session on the file system after you have corrected the cause of the failure.

The table below lists the supported object names for your operating system.

Table 62 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile views) • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Solaris	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Linux	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c • Linux MPIO's mpath device
AIX ^a	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device

- If the object name is on a Logical Volume Manager, then admsnap will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.
- If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

-t *object_type*

Specifies the object type. This switch is valid on Windows and UNIX operating systems only. Valid Windows object types are device and drive letter. Valid Solaris object types are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

You should use the `-t` option when the argument for the `-o` option is not in a standard and recognizable form for your operating system's object or device.

If you do not specify the `-t` option, `admsnap` will attempt to determine the correct object type and report a failure if it cannot.

`-c`

Specifies a consistent session. The session will preserve the point-in-time copy across a set of source LUNs. The SnapView driver will delay any I/O requests to the set of source LUNs until the session has started on all LUNs (thus preserving the point-in-time on the entire set of LUNs). Once you start a consistent session, you cannot add another source LUN to the session. In the event of a failure, the software will not start the session on any source LUN and will display an error message.

Note

When starting a consistent session, multiple object names must be separated by a comma with no spaces. If desired, you can specify different object names. For examples, refer to SnapView examples in *EMC VNX for Block Concepts*.

You would use consistent mode if:

- you cannot stop I/O to the source LUN(s) before starting a session.
- you want to prevent other LUNs from being added to the session. If desired, you can start a consistent session on a single source LUN to prevent other LUNs from being added to the session.
- the set of source LUNs that you will use to start the session spans both SPs; however, the source LUNs cannot span systems.

The consistent feature is available on a per-session basis (not per snapshot or source LUN) and counts as one of the eight sessions per source LUN limit.

Note

While a consistent session is in the process of starting on all LUNs, you cannot stop the session (the software may stop the session immediately after the start completes if I/O is queued and there is not enough space in the SP's reserved LUN pool), perform a rollback operation, or activate a snapshot to the session.

`-p`

Specifies a persistent session. If you use this option, the session will run in persistence mode and will survive SP failures, software installation upgrades, and trespass events. With `admsnap` version 2.24 or later, all SnapView sessions run in persistence mode, so the `-p` switch is no longer required.

EXAMPLE

Windows examples:

```
admsnap start -s session1 -o \\.\PhysicalDrive1
```

```
admsnap start -s session1 -o \\.\
PhysicalDrive1,PhysicalDrive2,PhysicalDrive3 -c
```

```
admsnap start -s session1 -o F:,G:,H:,I -c -p
```

```
admsnap start -s session1 -o F:,G:,\\.\PhysicalDrive1,PhysicalDrive2
-c UNIX examples:
```

```
admsnap start -s session1 -o /mnt/fs1 -p
```

```
admsnap start -s session1 -o /mnt/fs1,/mnt/fs2,/mnt/fs3 -c
```

```
admsnap start -s session1 -o /dev/sda,/dev/sdc,/dev/sdd -c -p
```

```
admsnap start -s session1 -o /dev/sda,/mnt/fs2,vol_name -c
```

admsnap stop

Stops the point-in-time copy of a source LUN.

PREREQUISITES

You must have administrative privileges on the production server.

DESCRIPTION

The `admsnap stop` operation stops a SnapView session. From the production server that owns the source LUN, the `stop` command stops a SnapView session, freeing the reserved LUN and disk space the session used.

The `admsnap stop` command is valid on Windows, Solaris, HP-UX, Linux, NetWare, and AIX systems.

Note

If you started a session on multiple source LUNs, you can select any of the source LUNs to stop the session.

SYNTAX

```
admsnap stop [-h] -s session-name -o object_name
[-t object_type]
```

OPTIONS

[-h]

Lists the help output for a given operation. If you use this option on the command line, a help message is displayed for each option you specify. `Admsnap` will not perform the command; it lists the help output only.

This option is valid for Windows, Solaris, HP-UX, Linux, and AIX systems.

-s *session_name*

Specifies the session name that you assigned when you started the session. You must specify the session name you used when you started the session. This is a string of up to 256 characters, and can be composed of letters, numbers and underscores.

The session name is case sensitive and must be unique on the system. If the session name contains spaces, you must enter the name in double quotation marks (for example, "June 21st Session").

-o *object_name*

Specifies an object name. You must specify an object name to identify the session.

If you use the `-t` option, the device name must be a valid name for the type of object specified with the `-t` option.

The table below lists the supported object names for your operating system.

Table 63 Operating system object names

Operating system	Supported object name(s)
Windows	<ul style="list-style-type: none"> • Drive letter • Physical drive name
HP-UX ^a	<ul style="list-style-type: none"> • Device name (Legacy and agile views) • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Solaris	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c
Linux	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device • Veritas volume name (disk-based and enclosure-based)^c • Linux MPIO's mpath device
AIX ^a	<ul style="list-style-type: none"> • Device name • File system name^b • Powerpath device

- If the object name is on a Logical Volume Manager, then admsnap will parse the Logical Volume Manager (LVM) to get the device names and will attempt to flush those devices.
- If the object name is a file system, then admsnap will parse the file system table entry to get the device name and will attempt to flush that device.
- If the object name is a Veritas volume name, admsnap will parse the Veritas volume to determine all physical disks for the volume and will attempt to flush each physical disk.

The only time admsnap will not stop the session on all devices within an object is when the device is also in another object that already has the same session name started.

`[-t object_type]`

Specifies the object type. This switch is valid on Windows or UNIX operating systems only. Valid types on Windows are device and drive letter. Valid types on Solaris are device, volume, and file system. If the object type name contains spaces, enter the name without spaces.

You should use the `-t` option when the argument for the `-o` option is not in a standard and recognizable form for your operating system's object or device. If you do not specify the `-t` option, admsnap will attempt to determine the correct object type and report a failure if it cannot.

EXAMPLE

Windows example:

```
admsnap stop -s session1 -o \\.\PhysicalDrive1
```


CHAPTER 13

EMC VNX Snapshots Software

This chapter introduces the EMC VNX Snapshots™ software and the associated commands. The VNX Snapshots commands are divided into two types: Snapshot CLI commands and SnapCLI commands.

Major topics include:

- [VNX Snapshots CLI commands](#)..... 602
- [SnapCLI commands for VNX Snapshots](#)..... 623

VNX Snapshots CLI commands

EMC VNX Snapshots is a VNX software feature. It allows creating point-in-time data copies similar to SnapView Snapshots. VNX Snapshots are not clones, and therefore do not necessarily take a lot of pool capacity. They take only seconds to create. Depending on your application needs, you can create VNX Snapshots and snapshots of VNX Snapshots. VNX Snapshots do not replace the existing SnapView technology, and these features can coexist. SnapView clones can have their own VNX Snapshots. VNX Snapshots can be used for data backups, software development and testing, re-purposing, data validation, and local rapid restores. VNX Snapshots technology writes new data to a new area within a pool, without the need to read and write the old data block. Thus improving the overall performance compared to SnapView Snapshot technology.

VNX Snapshots mount points are created by using Unisphere Secure CLI. You can use the `-allowInbandSnapAttach yes` option if you want to attach VNX Snapshots to the mount points, using SnapCLI.

VNX Snapshots supports 256 writable snaps per pool LUN. Branching, also known as snap of a snap, is also supported. There are no restrictions on the number of branches; however, the total number of Snapshots for a given primary LUN is within 256 members. Consistency groups (CG) are introduced as well. Several pool LUNs can be combined into a CG and snapped concurrently.

This section lists the Snapshot CLI commands that support the VNX Snapshots feature.

snap -feature

Lists the properties of the VNX Snapshots feature.

DESCRIPTION

The `snap -feature` command allows the user to list properties of VNX Snapshots feature.

SYNTAX

```
snap -feature -info
```

OUTPUT

```
naviseccli snap -feature -info
```

```
Is Advanced Snaps Supported:  True
Max. Snaps: 4096
Max. Snaps Per Source: 256
Max. CGs: 120
Max. Members Per CG: 64
Max. SnapLUNs: 512
Total Number of Snaps: 30
Total Number of CGs: 1
Total Number of SnapLUNs: 10
```

snap -create

Creates snapshots on thick LUNs or thin LUNs and Snapshot mount point.

DESCRIPTION

The `snap -create` command allows the user to create Snapshot on thick LUNs or thin LUNs and snapshot mount. It also allows creation of snapshot set on consistency group.

SYNTAX

```

snap -create -res resource
[-resType type] [-name snapName] [-descr description]
[-keepFor number {h|d|m|y}]
[-allowAutoDelete{yes|no}] [-allowReadWrite {yes|no}]
[-ignoreMigrationCheck] [-ignoreDeduplicationCheck]

```

OPTIONS

-res *resource*

Specifies the ID of the resource to be snapped.

-resType *type*

Specifies the type of the resource to be snapped. The valid values are **CG** (consistency group) and **LUN**. If this is not specified, `-create` checks whether a consistency group exists with the same name as the value specified in the `-res` option. If it does, the matching consistency group is snapped. If no consistency group exists, it will check if any LUN exists with its number same as the value specified in the `-res` switch. If LUN exists, it will be snapped. Otherwise, an error is returned. To create a Snapshot, you must use the LUN ID, not the LUN name.

Note

You must flush the host buffers before creating a Snapshot.

-name *snapName*

Specifies the name of the snapshot.

-descr *description*

Specifies the description of the snapshot.

-keepFor *number*{h|d|m|y}

Specifies the retention time after which the snapshots will be deleted. The interval can be defined in years, months, days, or hours. The format of the value *number*{h|d|m|y}

where:

number - An integer value.

h|d|m|y --The qualifier of the number. The valid values (case-insensitive) are h (hours), d (days), m (months) and y (years).

- If the qualifier is h (hours), the valid range is from 1 to 87600.
- If the qualifier is d (days), the valid range is from 1 to 3650.
- If the qualifier is m (months), the valid range is from 1 to 120.
- If the qualifier is y (years), the valid range is from 1 to 10.

-allowAutoDelete {yes|no}

Specifies whether auto delete is allowed on this Snapshot. Valid values are:

- Yes – allow auto delete
- No – disallow auto delete

The options `-keepFor` and `-allowAutoDelete` are mutually exclusive. If any `-keepFor` value is specified, the `-allowAutoDelete` is disabled.

-allowReadWrite {yes|no}

Specifies whether a user can attach one or more Snapshot mount points to this snapshot. The valid values are:

- Yes - attachable

- No - non-attachable

`-ignoreMigrationCheck`

Executes the command without prompting for confirmation.

`-ignoreDeduplicationCheck`

Overrides the error during snapshot creation while the source is enabling deduplication or disabling deduplication.

EXAMPLE #1

```
naviseccli snap -create -res 1 -name TLU1_Snapshot1 -keepFor 10d
-allowReadWrite yes
```

EXAMPLE #2

```
naviseccli -h riker-spa snap -create -res 1 -ignoreDeduplicationCheck
```

OUTPUT

There is no output for this command.

snap -destroy

Deletes a snapshot.

DESCRIPTION

The `snap -destroy` command allows the user to delete a snapshot.

SYNTAX

```
snap -destroy -id snapName [-o]
```

OPTIONS

`-id`

Specifies the ID of the snapshot to be deleted.

EXAMPLE

```
naviseccli snap -destroy -id TLU1_Snapshot1
```

OUTPUT

There is no output for this command.

snap -list

Lists the properties of selected snapshot or snapshot set.

DESCRIPTION

The `snap -list` command lists the properties of the specified snapshot. If no parameter is specified, all the basic properties of the specified snapshot are returned. If none of the snapshots matches the identifier parameters, nothing is returned.

SYNTAX

```
snap -list [{-id snapName |[-resType type] [ -res resource]]]
[{-brief|-detail}]
```

OPTIONS

`-id`

Specifies the ID of snapshot or snapshot set to be listed.

`-resType`

Specifies the type of snapshot or snapshot set to be listed. The valid values are:

- lun - list snapshot

- `cg - list snapshot set`

`-brief`

Shows only the basic properties of a snapshot or snapshot set.

`-detail`

Shows all the properties of a snapshot or snapshot set.

Note

If neither `-brief` nor `-detail` option is specified, the basic properties are shown.

EXAMPLE #1

To show all the snapshots on the array:

`naviseccli snap -list`

```
Name: LUN0_2009-04-13-11:13:31_2
Description: Snapshot before update to new exchange server
Creation time: 04/13/09 11:13:31
Source LUNs: 0
Source CG: N/A
State: Ready
Allow Read/Write: Yes
Modified: Yes
Allow auto delete: Yes
Expiration date: Never

Name: LUN1_2009-04-15-09:01:01_8
Description: Snapshot before installing Visual Studio
Creation time: 04/15/09 09:01:01
Source LUNs: 1
Source CG: N/A
State: Ready
Allow Read/Write: No
Modified: No
Allow auto delete: No
Expiration date: 04/25/09 09:01:01

Name: CG0_2009-04-13-11:30:31_2
Description: Snapshot of CG0 before update to new sql server
Creation time: 04/13/09 11:30:31
Source LUNs: 11, 12
Source CG: CG0
State: Ready
Allow Read/Write: Yes
Modified: No
Allow auto delete: Yes
Expiration date: Never
```

EXAMPLE #2

To show the snapshot with the specified ID with detailed output:

`naviseccli snap -list -id CG0_2009-04-13-11:30:31_2 -detail`

```
Name: CG0_2009-04-13-11:30:31_2
Description: Snapshot of CG0 before update to new sql server
Creation time: 04/13/09 11:30:31 Last modify time: N/A
Last modified by: N/A
Source LUNs: 11, 12
Source CG: CG0 Primary LUNs: 11, 12
State: Ready
Status: N/A
Allow Read/Write: Yes
Modified: No
Attached LUN(s): 4091,4092
```

```
Allow auto delete: Yes
Expiration date: Never
```

EXAMPLE #3

To show all the snapshots of the specified resource type:

```
naviseccli snap -list -resType lun
```

```
Name: LUN0_2009-04-13-11:13:31_2
Description: Snapshot before update to new exchange server
Creation time: 04/13/09 11:13:31
Source LUNs: 0
Source CG: N/A
State: Ready
Allow Read/Write: Yes Modified: Yes
Allow auto delete: Yes
Expiration date: Never
```

```
Name: LUN1_2009-04-15-09:01:01_8
Description: Snapshot before installing Visual Studio
Creation time: 04/15/09 09:01:01
Source LUNs: 1
Source CG: N/A State: Ready
Allow Read/Write: No
Modified: No
Allow auto delete: No
Expiration date: 04/25/09 09:01:01
```

EXAMPLE #4

To show all the snapshots related to the specified source resource, with detailed output:

```
naviseccli snap -list -res 0 -detail
```

```
Name: LUN0_2009-04-13-11:13:31_2
Description: Snapshot before update to new exchange server
Creation time: 04/13/09 11:13:31
Last modify time: 04/15/09 12:30:20
Last modified by: SnapLUN4
Source LUNs: 0
Source CG: N/A
Primary LUNs: 0
State: Ready
Status: N/A
Allow Read/Write: Yes
Modified: Yes
Attached LUN(s): N/A
Allow auto delete: Yes
Expiration date: Never
```

snap -modify

Modifies the snapshot's name, description, expiration time, auto delete policy, or attach policy.

DESCRIPTION

The `snap -modify` command allows the user to give a different name, description, expiration time, auto delete policy, or attach policy to the Snapshot.

SYNTAX

```
snap -modify -id snapName [-name newName] [-descr description]
[{-keepFor number {h|d|m|y}}
|-allowAutoDelete {yes|no}}]
[-allowReadWrite{yes|no}]
```

OPTIONS

- id *snapName*
Specifies the ID of Snapshot to be modified.
- name *newName*
Specifies the new name for the Snapshot. This option changes the ID as well.
- descr *description*
Specifies the new description for the Snapshot.
- keepFor *number*{h|d|m|y}
Specifies the retention time after which the snapshots will be deleted. The interval can be defined in years, months, days, or hours. The format of the value is: *number* {h|d|m|y}.
where:
number - An integer value.
{h|d|m|y} - The qualifier of the number. The valid values (case-insensitive) are h (hours), d (days), m (months) and y (years).
 - If the qualifier is h (hours), the valid range is from 1 to 87600.
 - If the qualifier is d (days), the valid range is from 1 to 3650.
 - If the qualifier is m (months), the valid range is from 1 to 120.
 - If the qualifier is y (years), the valid range is from 1 to 10.
- allowAutoDelete {yes|no}
Specifies the new auto delete policy. The valid values are:
 - Yes - allow auto delete
 - No - disallow auto delete
- allowReadWrite {yes|no}
Specifies whether a user can attach Snapshot mount point(s) to this Snapshot. The valid values are:
 - Yes - attachable
 - No - non-attachable

EXAMPLE

```
naviseccli snap -modify -id TLU1_Snapshot1 -name TLU1_Snapshot2
-allowAutoDelete yes
```

OUTPUT

There is no output for this command.

snap -copy

Replicates a Snapshot.

DESCRIPTION

The `snap -copy` command allows the user to replicate a Snapshot.

SYNTAX

```
snap -copy -id snapName [-name newName] [-ignoreMigrationCheck]
[-ignoreDeduplicationCheck]
```

OPTIONS

- id
Specifies the ID of a Snapshot to be replicated.

-name
Specifies the new Snapshot name.

-ignoreMigrationCheck
Executes the command without prompting for confirmation.

-ignoreDeduplicationCheck
Overrides the error during snapshot creation while the source is enabling deduplication or disabling deduplication.

EXAMPLE

```
naviseccli -h riker-spa snap -copy -id snap1 -ignoreDeduplicationCheck
```

OUTPUT

```
naviseccli snap -copy -id TLU1_Snapshot1 -name TLU1_Snapshot2
ID = TLU1_Snapshot2
Operation completed successfully
```

snap -restore

Restores a primary LUN, snapshot mount point, or consistency group.

DESCRIPTION

The `snap -restore` command allows a primary LUN, snapshot mount point, or consistency group to be rolled back to a previous point in time, which is defined by a snapshot.

SYNTAX

```
snap -restore -id snapName [-bakName bakName] [-res lunNumber]
```

OPTIONS

-id
Specifies the ID of the snapshot to be restored.

-bakName
Specifies the backup snapshot name.

Note

To save the current file, a backup snapshot is created on the primary LUN before restoring.

-res
Specifies the LUN number of the primary LUN, snapshot mount point, or one of the LUNs in snapshot set to be restored.

Note

- Must be specified when restoring the Snapshot
- Optional when restoring the Snapshot Set

EXAMPLE

```
naviseccli snap -restore -id TLU1_Snapshot1 -bakName TLU1_Snapshot2
-res 1
```

OUTPUT

There is no output for this command.

snap -attach

Attaches snapshot mount points to the snapshot.

DESCRIPTION

The `snap -attach` command allows the user to attach snapshot mount point(s) to the snapshot, so that the data of the Snapshot is exposed to host.

SYNTAX

```
snap -attach -id snapName -res lunNumber
```

OPTIONS

`-id`

Specifies the ID of snapshot to be attached.

`-res`

Specifies the LUN number of snapshot mount point to be attached.

EXAMPLE

```
naviseccli snap -attach -id TLU1_Snapshot1 -res 4
```

OUTPUT

There is no output for this command.

snap -detach

Detaches snapshot mount points from a snapshot.

DESCRIPTION

The `snap -detach` command allows the user to detach snapshot mount points from a snapshot.

SYNTAX

```
snap -detach -id snapName [-res lunNumber] [-o]
```

OPTIONS

`-id`

Specifies the ID of snapshot to be detached from.

`-res`

Specifies the LUN number of snapshot mount point to be detached.

`-o`

Overrides the warning on detaching a snapshot mount point that has snapshots.

EXAMPLE

```
naviseccli snap -detach -id TLU1_Snapshot1 -res 4
```

OUTPUT

There is no output for this command.

snap -group -create

Creates a consistency group.

DESCRIPTION

The `snap -group -create` command allows user to create a consistency group.

SYNTAX

```
snap -group -create -name cgName [-res lunNumber(s)] [-descr description] [-allowSnapAutoDelete {yes|no} ]
```

OPTIONS

- name
Specifies the name of the consistency group to be created.
- res
Specifies the LUN number of the LUNs to be added into the consistency group. If there are multiple LUN IDs, then they should be separated by commas.
- descr
Specifies the description of the consistency group.
- allowSnapAutoDelete
Specifies the default auto delete policy, which will apply on snapshot set of this CG. The valid values are:
 - *yes* - allow auto delete
 - *no* - disallow auto delete

EXAMPLE

```
naviseccli snap -group -create -name CG1_Exchange -res 2,3
```

OUTPUT

There is no output for this command.

snap -group -destroy

Deletes a consistency group.

DESCRIPTION

The `snap -group -destroy` command allows user to delete a consistency group.

SYNTAX

```
snap -group -destroy -id cgName [-destroySnapshots]
```

OPTIONS

- id
Specifies the ID of the consistency group to be deleted.
- destroySnapshots
Forces all the snapshot set associated with this consistency group to be deleted.

EXAMPLE

```
naviseccli snap -group -destroy -id CG1_Exchange
```

OUTPUT

There is no output for this command.

snap -group -modify

Modifies a consistency group name, description, or auto delete policy.

DESCRIPTION

The `snap -group -modify` command allows user to modify consistency group name, description, or auto delete policy.

SYNTAX

```
snap -group -modify -id cgName [-name newName] [-descr description] [-allowSnapAutoDelete {yes|no}]
```

OPTIONS

- id
Specifies the ID of the Consistency Group to be modified.

- name
Specifies a new name for the Consistency Group. The ID will also be changed.
- descr
Specifies a new description for the Consistency Group
- allowSnapAutoDelete
Specifies the default auto delete policy that will apply to the Snapshot Set of this Consistency Group. Valid values are:
 - *yes* - allow auto delete
 - *no* - disallow auto delete

EXAMPLE

```
naviseccli snap -group -modify -id CG1_Exchange -allowSnapAutoDelete
yes
```

OUTPUT

There is no output for this command.

snap -group -list

Lists the properties of the selected or all consistency groups.

DESCRIPTION

The `snap -group -list` command allows the user to list properties of selected or all consistency groups. If none of the consistency group matches the identifier parameters passed, nothing will be returned.

SYNTAX

```
snap -group -list -id cgName [{-brief|-detail}]
```

OPTIONS

- id
Specifies the ID of consistency groups to be listed.
- brief
Show only the basic properties of a consistency group.
- detail
Show all properties of a consistency group.

EXAMPLE #1

To show all consistency groups on the array:

```
naviseccli snap -group -list
```

```
Name: CG123
Description: Consistency Group containing all LUNs used in exchange
Allow auto delete: Yes
Member(s): 1, 3, 5, 6, 10
State: Ready

Name: CG_SQLServer
Description: Consistency Group containing all LUNs used in SQL Server
Allow auto delete: Yes
Member(s): 2, 4
State: Ready
```

EXAMPLE #2

To show consistency groups with the specified ID, with detail output:

```
naviseccli snap -group -list -id CG_SQLServer -detail
```

```
Name: CG_SQLServer
Description: Consistency Group containing all LUNs used in SQL Server
Allow auto delete: Yes
Member(s): 2, 4
State: Ready
Status: N/A
```

snap -group -addmember

Adds member LUNs to the consistency group.

DESCRIPTION

The `snap -group -addmember` command adds member LUNs to the consistency group.

SYNTAX

```
snap -group -addmember -id cgName -res lunNumber(s)
```

OPTIONS

`-id`

Specifies the ID of consistency group to which the LUNs will be added.

`-res`

Specifies ID of LUNs to be added to the consistency group. Multiple LUN IDs should be separated by commas.

EXAMPLE

```
naviseccli snap -group -addmember -id CG1_Exchange -res 2
```

OUTPUT

There is no output for this command.

snap -group -rmmember

Removes member LUN from the consistency groups.

DESCRIPTION

The `snap -group -rmmember` command removes member LUN from the consistency groups.

SYNTAX

```
snap -group -rmmember -id cgName -res lunNumber(s)
```

OPTIONS

`-id`

Specifies the ID of consistency group from which the LUN will be removed.

`-res`

Specifies ID of LUNs to be removed from the consistency groups. Multiple LUN IDs should be separated by commas.

EXAMPLE

```
naviseccli snap -group -rmmember -id CG1_Exchange -res 2
```

OUTPUT

There is no output for this command.

snap -group -replmember

Replaces member LUNs of the consistency groups with the specified LUNs.

DESCRIPTION

The `snap -group -replmember` command replaces member LUNs of the consistency groups with the specified LUNs.

SYNTAX

```
snap -group -replmember -id cgName -res lunNumber(s)
```

OPTIONS

`-id`
Specifies the ID of the consistency group of which the LUNs will be replaced.

`-res`
Specifies ID of LUNs to replace the existing members in the consistency group.

EXAMPLE

```
naviseccli snap -group -replmember -id CG1_Exchange -res 2,3
```

OUTPUT

There is no output for this command.

lun -create

Creates snapshot mount point to allow the user to attach a snapshot on it, so that the snapshot data can be exposed to the host.

DESCRIPTION

The `lun -create` command creates snapshot mount point to allow user to attach a snapshot on it, so that the snapshot data can be exposed to the host.

SYNTAX

```
lun -create -type Snap -primaryLun primaryLunNumber |  
-primaryLunName primaryLunName [-l lunNumber] [-name lunName]  
[-allowSnapAutoDelete yes|no] [-allowInbandSnapAttach yes|no]  
[-sp spID]
```

OPTIONS

`-type`
Specifies the type of the new MLU; the value must be `Snap` to create snapshot mount point.

`-primaryLun`
Specifies the ID of the primary LUN.

`-primaryLunName`
Specifies the name of the primary LUN.

`-l`
Specifies the snapshot mount point number.

`-name`
Specifies the snapshot mount point name.

`-allowSnapAutoDelete`
Specifies the snapshot auto delete policy of the snapshot mount point. Valid values are:

- Yes - allow auto delete

- No - disallow auto delete
- `-allowInbandSnapAttach`
Specifies whether to allow host-based SnapCLI to attach snapshots. Valid values are:
- Yes - allow
 - No - disallow
- `-sp`
Specifies the owner SP of the snapshot mount point to be created.

OUTPUT

There is no output for this command.

lun -create

Creates a LUN by specifying an auto delete policy.

DESCRIPTION

The `lun -create` command allows the user to create a LUN by specifying an auto delete policy.

SYNTAX

```
lun -create
[-type Thin|NonThin] -capacity capacity [-sq mb|gb|tb|bc]
-poolId storagePoolID|-poolName storagePoolName -sp spID
[-aa autoAssignment] [-l lunNumber] [-name lunName]
[-offset logicalBlockAddress] [-initialTier optimizePool|
highestAvailable|lowestAvailable] [-allowSnapAutoDelete yes|no]
[-ignoreThresholds] [-allocationPolicy onDemand|automatic]
```

OPTIONS

- `-offset`
Specifies the LUN to be offset.
- `-allowSnapAutoDelete`
Specifies the snapshot auto delete policy of the LUN to be created. Valid values are:
- *yes* - allow auto delete
 - *no* - disallow auto delete
- `-allocationPolicy`
Specifies the allocation policy on the pool LUN. Valid values are:
- *onDemand*
 - *automatic* (default)
- `-ignoreThresholds`
Forces the non-snap LUN to be created, ignoring possible threshold related error.

OUTPUT

There is no output for this command.

lun -destroy

Deletes the primary LUN or snapshot mount point.

DESCRIPTION

The `lun -destroy` command allows the user to delete the primary LUN or snapshot mount point.

SYNTAX

```
lun -destroy -l lunNumber -name lunName [-destroySnapshots]
[-forceDetach] [-o]
```

OPTIONS

-l
Specifies the number of the LUN to be deleted.

-name
Specifies the name of the LUN to be deleted.

-destroySnapshots
Forces deletion of the primary LUN along with all of its snapshots.

-forceDetach
Forces deletion of the snapshot mount point even if it is currently attached.

-o
Overrides the confirmation message.

OUTPUT

There is no output for this command.

lun -modify

Modifies the primary LUN or snapshot mount point name and auto delete policy.

DESCRIPTION

The `lun -modify` command modifies the primary LUN or snapshot mount point name, auto delete policy, and allows inband snap attach policy.

SYNTAX

```
lun -modify -l lunNumber|-name lunName [-aa autoAssign]
[-newName newLUNName] [-sp newDefaultSPID]
[-allowSnapAutoDelete yes|no] [-allowInbandSnapAttach yes|no]
[-allocationPolicy automatic]
[-deduplication On|Off [-destroySnapshots]] [-o]
```

OPTIONS

-l
Specifies the number of LUNs to be modified.

-name
Specifies the name of LUN to be modified.

-aa
Specifies the auto assignment.

-newName
Specifies the new name for the specified LUN.

-sp
Specifies the new default SP of the LUN to be modified.

-allowSnapAutoDelete
Specifies the snapshot auto delete policy of the LUN to be modified. The valid values are:

- *yes* - allow auto delete
- *no* - disallow auto delete

-allowInbandSnapAttach
Specifies whether the host-based SnapCLI can attach snapshots. Valid values are:

- *yes* - allow

- *no* - disallow
- `-allocationPolicy`
Specifies the allocation policy on the pool LUN. The only valid value is automatic.
- `-o`
Overrides the confirmation message.
- `-deduplication On|Off`
Specifies whether deduplication is enabled on this LUN. Valid values are On or Off.
- `-destroySnapshots`
Specifies to explicitly destroy associated VNX Snapshots before turning deduplication on or off for the specified LUN. If a LUN has associated VNX Snapshots and you do not use the `-destroySnapshots` option, you cannot change the deduplication value.

OUTPUT

There is no output for this command.

lun -expand

Expands the primary LUN or snapshot mount point.

DESCRIPTION

The `lun -expand` command expands the primary LUN or snapshot mount point.

SYNTAX

```
lun -expand -l lunNumber -name lunName -capacity capacity
[-sq mb|gb|tb|bc] [-ignoreThresholds] [-o]
```

OPTIONS

- `-ignoreThresholds`
Force the non-snap LUN to be created, ignoring the possible threshold-related error.

OUTPUT

There is no output for this command.

lun -attach

Attaches the snapshot mount point to the specified snapshot.

DESCRIPTION

The `lun -attach` command allows the user to attach snapshot mount point to the specified snapshot.

SYNTAX

```
lun -attach
-l lunNumber -name lunName -snapName snapName
```

OPTIONS

- `-l`
Specifies the LUN number to be attached.
- `-name`
Specifies the name of the LUN to be detached.
- `-snapName`
Specifies the snapshot name to be attached to.

OUTPUT

There is no output for this command.

lun -detach

Detaches the snapshot mount point from the attached snapshot.

DESCRIPTION

The `lun -detach` command allows the user to detach the snapshot mount point from the attached snapshot.

SYNTAX

```
naviseccli lun -detach -l lunNumber [-name lunName]
```

OPTIONS

- l
Specifies the LUN number to be detached.
- name
Specifies the name of the LUN to be detached.

OUTPUT

There is no output for this command.

storagepool -create

Creates a storage pool by specifying a snapshot auto delete policy and MR-related thresholds.

DESCRIPTION

The `storagepool -create` command allows the user to create a storage pool by specifying a snapshot auto delete policy and MR-related thresholds.

SYNTAX

```
storagepool -create [-snapPoolFullThresholdEnabled on|off]
[-snapPoolFullHWM poolFullHWM] [-snapPoolFullLWM poolFullLWM]
[-snapSpaceUsedThresholdEnabled on|off]
[-snapSpaceUsedHWM snapSpaceUsedHWM]
[-snapSpaceUsedLWM snapSpaceUsedLWM]
```

OPTIONS

- snapPoolFullThresholdEnabled
Indicates that the system should check the pool full high water mark for auto delete. Valid values are:
 - *on* - enabled
 - *off* - disabled
- snapPoolFullHWM
Specifies the pool full high watermark (that will trigger the auto delete process) of the storage pool to be created.
- snapPoolFullLWM
Specifies the pool full low watermark (that will stop the auto delete process) of the storage pool to be created.
- snapSpaceUsedThresholdEnabled
Indicates that the system should check the snapshot space used high water mark for auto delete. Valid values are:
 - *on* - enabled
 - *off* - disabled
- snapSpaceUsedHWM

Specifies the snapshot space used high water mark that will trigger the auto delete process of the storage pool to be created.

`-snapSpaceUsedLWM`

Specifies the snapshot space used low water mark that will stop the auto delete process of the storage p to be created.

OUTPUT

There is no output for this command.

storagepool -modify

Modifies the snapshot auto delete policy and MR-related thresholds of an existing storage pool.

DESCRIPTION

The `storagepool -modify` command allows the user to modify the snapshot auto delete policy and MR-related thresholds of an existing storage pool.

SYNTAX

```
storagepool -modify [-snapPoolFullThresholdEnabled on|off]
[-snapPoolFullHWM poolFullHWM] [-snapPoolFullLWM poolFullLWM]
[-snapSpaceUsedThresholdEnabled on|off]
[-snapSpaceUsedHWM snapSpaceUsedHWM]
[-snapSpaceUsedLWM snapSpaceUsedLWM]
```

OPTIONS

`-snapPoolFullThresholdEnabled`

Indicates that the system should check the pool full high water mark for auto delete.

Valid values are:

- *on* - enabled
- *off* - disabled

`-snapPoolFullHWM`

Specifies the new pool full high water mark (that will trigger the auto delete process) of the Storage Pool to be modified.

`-snapPoolFullLWM`

Specifies the new pool full low water mark (that will stop the auto delete process) of the Storage Pool to be modified.

`-snapSpaceUsedThresholdEnabled`

Indicates the system to check the snapshot space used high water mark for auto delete. Valid values are:

- *on* - enabled
- *off* - disabled

`-snapSpaceUsedHWM`

Specifies the new snapshot space used high water mark (that will trigger the auto delete process) of the Storage Pool to be modified.

`-snapSpaceUsedLWM`

Specifies the new snapshot space used low water mark (that will stop the auto delete process) of the Storage Pool to be modified.

OUTPUT

There is no output for this command.

storagepool -list

Lists the total snapshot capacity, auto delete policy, auto delete state, and MR related thresholds of the specified storage pool or pools.

DESCRIPTION

The `storagepool -list` command allows the user to list the total snapshot capacity, auto delete policy, auto delete state, and MR-related thresholds of specified storage pool or pools.

SYNTAX

```
storagepool -list
[-lunAllocation] [-snapshotAllocation]
[-metadataAllocation] [-lunSubscribedCap]
[-snapshotSubscribedCap] [-metadataSubscribedCap]
[-compressionSavings] [-capacities]
[-snapPoolFullThresholdEnabled] [-snapPoolFullHWM]
[-snapPoolFullLWM] [-snapPoolFullState]
[-snapSpaceUsedThresholdEnabled] [-snapSpaceUsedHWM]
[-snapSpaceUsedLWM] [-snapSpaceUsedState]
```

OPTIONS

`-lunAllocation`

The total amount of storage that has been allocated from the pool for data written to all the primary LUNs in the pool.

`-snapshotAllocation`

The total amount of storage that has been allocated from the pool for all the VNX snapshots created in the pool.

`-metadataAllocation`

The total amount of storage that has been allocated from the pool to store metadata for the LUNs and VNX snapshots in the pool.

`-lunSubscribedCap`

The total amount of user data that can be written to all of the LUNs in the pool.

`-snapshotSubscribedCap`

The total amount of space that all LUN snapshots currently consume as well as what they would consume if all of the primary LUNs, data was overwritten.

`-metadataSubscribedCap`

The total amount of metadata storage that is required to support writing to the entire contents of all of the LUNs in the pool.

`-compressionSavings`

The total amount of storage that is being saved by using compression on the LUNs in the pool.

`-capacities`

Shows all related capacities including “Raw Capacity”, “User Capacity”, “Consumed Capacity”, “LUN Allocation”, “Snapshot Allocation”, “Metadata Allocation”, “Available Capacity”, “Percent Full”, “Total Subscribed Capacity”, “LUN Subscribed Capacity”, “Snapshot Subscribed Capacity”, “Metadata Subscribed Capacity”, “Percent Subscribed”, “Oversubscribed by”, and “Compression Savings”.

`-snapPoolFullThresholdEnabled`

Shows whether the system should check the pool full high water mark for auto delete.

`-snapPoolFullHWM`

Shows the pool full high water mark of the specified storage pool or pools.

`-snapPoolFullLWM`

Shows the pool full low water mark of the specified storage pool or pools.

- snapPoolFullState
Shows the pool full auto delete state.
- snapSpaceUsedThresholdEnabled
Shows whether the system should check the snapshot space used high water mark for auto-delete.
- snapSpaceUsedHWM
Shows snapshot space used high water mark of specified storage pool or pools.
- snapSpaceUsedLWM
Shows snapshot space used low water mark of specified storage pool or pools.
- snapSpaceUsedState
Shows snapshot space used auto delete state:
 - Idle
 - Running
 - Could not Reach LWM
 - System Paused
 - Failed

EXAMPLE

```
naviseccli storagepool -list -all
```

OUTPUT

```
Pool Name: Pool 0
Pool ID: 0
Raid Type: r_5
Percent Full Threshold: 70
Description:
Disk Type: SAS
State: Ready
Status: OK(0x0)
Current Operation: None
Current Operation State: N/A
Current Operation Status: N/A
Current Operation Percent Completed: 0
Raw Capacity (Blocks): 2814421510
Raw Capacity (GBs): 1342.021
User Capacity (Blocks): 2246234112
User Capacity (GBs): 1071.088
Consumed Capacity (Blocks): 39923712
Consumed Capacity (GBs): 19.037
LUN Allocation (Blocks): 22063104
LUN Allocation (GBs): 10.521
Snapshot Allocation (Blocks): 22063104
Snapshot Allocation (GBs): 10.521
Metadata Allocation (Blocks): 22063104
Metadata Allocation (GBs): 10.521
Available Capacity (Blocks): 2206310400
Available Capacity (GBs): 1052.051
Percent Full: 1.777
Total Subscribed Capacity (Blocks): 44126208
Total Subscribed Capacity (GBs): 21.041
LUN Subscribed Capacity (Blocks): 22063104
LUN Subscribed Capacity (GBs): 10.521
Snapshot Subscribed Capacity (Blocks): 22063104
Snapshot Subscribed Capacity (GBs): 10.521
Metadata Subscribed Capacity (Blocks): 22063104
Metadata Subscribed Capacity (GBs): 10.521
Percent Subscribed: 1.964
Oversubscribed by (Blocks): 0
Oversubscribed by (GBs): 0.000
Compression Savings (Blocks): N/A
Compression Savings (GBs): N/A
```

```

Tier Name: Performance
Raid Type: r_5
User Capacity (GBs): 1071.09
Consumed Capacity (GBs): 9.02
Available Capacity (GBs): 1062.07
Percent Subscribed: 0.84%
Disks (Type):
Bus 0 Enclosure 0 Disk 10 (SAS)
Bus 0 Enclosure 0 Disk 12 (SAS)
Bus 0 Enclosure 0 Disk 14 (SAS)
Bus 0 Enclosure 0 Disk 11 (SAS)
Bus 0 Enclosure 0 Disk 13 (SAS)

Disks:
Bus 0 Enclosure 0 Disk 14
Bus 0 Enclosure 0 Disk 12
Bus 0 Enclosure 0 Disk 10
Bus 0 Enclosure 0 Disk 13
Bus 0 Enclosure 0 Disk 11
LUNs: 200, 100, 101, 300
FAST Cache: N/A

Auto-Delete Pool Full Threshold Enabled: On
Auto-Delete Pool Full High Water Mark: 95
Auto-Delete Pool Full Low Water Mark: 85
Auto-Delete Pool Full State: Running
Auto-Delete Snapshot Space Used Threshold Enabled: On
Auto-Delete Snapshot Space Used High Water Mark: 25
Auto-Delete Snapshot Space Used Low Water Mark: 20
Auto-Delete Snapshot Space State: Idle

```

storagepool -resumeAutoDelete

Resumes auto-delete on an existing storage pool.

DESCRIPTION

The `storagepool -resumeAutoDelete` command allows the user to resume snapshots auto delete on an existing storage pool.

SYNTAX

```
storagepool -resumeAutoDelete -id poolID | -namepoolName
```

OPTIONS

`-id`

Specifies the ID of the storage pool to be operated on.

`-name`

Specifies the name of the storage pool to be operated on.

OUTPUT

There is no output for this command.

getlun

Gets the snapshot mount points as a LUN.

DESCRIPTION

The `getlun` command allows the user to get snapshot mount points as a LUN.

SYNTAX

getlun

```

-edit <lun-number> <-aa> <-at> <-bind> <-bp>
<-bread> <-brw> <-busy> <-bwrite> <-capacity> <-crus>
<-default> <-disk> <-drivetype> <-dsa> <-element> <-ena> <-ff>
<-hrdscsi> <-idle> <-idt> <-it> <-mp> <-name> <-nminus> <-owner>
<-parityelements> <-pd> <-pic> <-prb> <-prefetch> <-prf> <-ps>
<-psm> <-qavg> <-qmax> <-que> <-rb> <-rc> <-rd3> <-rchm>
<-reads> <-remap> <-ret> <-rh> <-rhist> <-rr> <-rwf> <-rwr>
<-service> <-sftscsi> <-srcp> <-ss> <-ssm> <-stripe> <-totque>
<-type> <-upb> <-verify> <-was> <-wc> <-wch> <-uid> <-wh>
<-whist> <-wr> <-writes> <-rg> <-state> <-private> <-usage> <-sc>
<-offset> <-addroffset> <-ismetalun> <-isthinlun> <-ispoollun>
<-idleticks> <-busyticks> <-lunidleticks> <-luncache>
<-lunbusyticks> <-trespass> <-status> <-brwsp> <-rwrsp>
<-lunbusytickssp> <-lunidletickssp> <-nanzq> <-sqlah> <-slst>
<-slct> <-etp> <-itp> <-nzcra> <-sor> <-wcrh> <-fwc> <-fastcache>
<-issnapmountpoint> <-all>

```

EXAMPLE

```
naviseccli getlun 1 -all
```

OUTPUT

```

Statistics logging is disabled.
Certain fields are not printed if statistics
logging is not enabled.

```

```

Prefetch size (blocks) =          0
Prefetch multiplier =           0
Segment size (blocks) =          0
Segment multiplier =            0
Maximum prefetch (blocks) =      0
Prefetch Disable Size (blocks) = 0
Prefetch idle count =           0

```

```

Prefetching: NO
Prefetched data retained      NO

```

```

Read cache configured according to
specified parameters.

```

```

Name                               LUN 1
Minimum latency reads N/A

RAID Type:                          N/A
RAIDGroup ID:                       N/A
State:                               Bound
Stripe Crossing:                    0
Element Size:                       0
Current owner:                      SP A
Offset:                              N/A
Auto-trespass:                      DISABLED
Auto-assign:                        DISABLED
Write cache:                         ENABLED
Read cache:                         ENABLED
Idle Threshold:                     0
Idle Delay Time:                    0
Write Aside Size:                   0
Default Owner:                      SP A
Rebuild Priority:                   N/A
Verify Priority:                    N/A
Prct Reads Forced Flushed:         0
Prct Writes Forced Flushed:        0
Prct Rebuilt:                       100

```

```

Prct Bound:                100
LUN Capacity (Megabytes):  10240
LUN Capacity (Blocks):    20971520
UID:
60:06:01:60:44:80:28:00:F6:E1:44:96:8C:D3:E0:11
LUN Capacity (Stripes):   N/A
Shrink State:              N/A
Is Private:                NO
Snapshots List:            Not Available
MirrorView Name if any:   Not Available
Address Offset:            N/A
Is Meta LUN:               NO
Is Thin LUN:               NO
Is Pool LUN:               YES
Is Snapshot Mount Point:  YES
LUN Idle Ticks:           N/A
LUN Busy Ticks:           N/A
LUN Offline (Cache Dirty Condition): N/A
LU Storage Groups:
Device Map:                 Valid
Average Read Time:         0
Average Write Time:        0
FAST Cache :                N/A
FAST Cache Read Hits:     N/A
FAST Cache Read Misses:   N/A
FAST Cache Write Hits:    N/A
FAST Cache Write Misses:  N/A

```

SnapCLI commands for VNX Snapshots

The snapCLI utility runs on host systems in conjunction with VNX Snapshot running on the VNX series SPs. It lets you create, copy, attach, detach, and destroy VNX Snapshots. All snapCLI commands are sent to the storage system through the Fibre Channel bus.

The snapCLI utility is an executable program that you can run interactively or with a script.

For all Microsoft Windows and UNIX platforms, flush the cached data from the production host before you create a snapshot. Also, flush the cached data on the secondary host if you have modified the VNX Snapshot. Failure to do this may result in an inconsistent or unusable point-in-time copy.

This section lists the snapCLI commands that support the VNX Snapshots feature.

snapcli create

Creates a snapshot of a pool-based LUN.

DESCRIPTION

The `snapcli create` command creates a snapshot under `snap_name` on `object_name` with optional `object_type`. You must specify a `snap_name` to identify the snapshot and an `object_name` to identify the members of the snapshot.

This operation is performed on the host that has access to the primary LUN that will be snapped. You must create a snapshot for each mount point of a specific primary LUN that you want to access.

SYNTAX

```

snapcli create -h | -s snap_name -o object_name [-t
object_type] -c cg_name

```

OPTIONS

-h

Prints out this message. If this option is present on the command line, no other option validation is performed.

-s

This option specifies a snapshot name. This is a string of up to 256 characters, and can be composed of letters, numbers, or underscores. The `snap_name` must be unique on the array.

-o

This option specifies an object-name. The object-name must be a valid object name that is one of the following object types:

- A device name
- A file system mount point
- A Veritas volume name
- A drive letter

If -c option is used, then a set of object names can be given on the command line with the object names separated by a comma field separator.

-t

This option specifies the object-type. The only valid object-types are "device", "filesystem", "volume", and "driveletter". This option is optional for these four supported object types. If not provided, SnapCLI attempts to determine the correct object type and reports a failure if it cannot.

-c

This option identifies a consistency group of LUNs, that could be snapped together, as follow:

- If the specified `cg_name` does not exist and a list of `object_names` is specified, the array attempts to create a new consistency group named `cg_name` for the `object_names` specified, and then snaps the consistency group. However, if any of the specified LUNs is in another consistency group, the create snap operation fails.
- If the `cg_name` already exists and the list of `object_names` is specified, its member LUNs are updated with the specified `object_names`, and the consistency group is snapped.

snapcli destroy

Destroys a named snapshot.

DESCRIPTION

This operation is used to destroy a snapshot named `snap_name` on `object_name` with optional `object_type`. You must specify a `snap_name` to identify the snapshot and an `object_name` to identify the members of the snapshot.

This operation is performed on the host that has access to the primary LUN that was previously snapped. Only one device per VNX storage array, currently in the snapshot, is needed to destroy the snapshot. A snapshot is not allowed to be destroyed if it is still attached to a mount point. If the named snap is for a consistency group, the consistency group will not be removed.

SYNTAX

```
snapcli destroy -h | -s snap_name -o object_name [-t object_type]
```

OPTIONS

-h

This option prints out this message. If this option is present on the command line, no other option validation is performed.

-s

This option specifies a snapshot name. This is a string of up to 256 characters, and can be composed of letters, numbers, or underscores. The *snap_name* must be unique on the array.

-o

This option specifies an *object_name*. The *object_name* must be a valid object name that is one of the following object types:

- A device name
- A file system mount point
- A Veritas volume name
- A drive letter

-t

This option specifies the *object_type*. The only valid *object_types* are "device", "filesystem", "volume", and "driveletter". This option is optional for these four basic supported object types. If no option is provided, SnapCLI attempts to determine the correct object type among these four object types, and reports a failure if it cannot.

snapcli copy

Creates a copy of the snapshot.

DESCRIPTION

You must specify the name of the snap to be created and the name of the snap to be snapped.

This operation is performed on the host that has access to the primary LUN or any member of the snap family. An attached snapshot cannot be copied, therefore the request to copy it will fail.

SYNTAX

```
snapcli copy -h | -s object_type -b snap_name -o object_name
```

OPTIONS

-h

Prints out this message. If this option is present on the command line, no other option validation is performed and this message is printed.

-s

Specifies a snapshot name for the snap to be created. This is a string of up to 256 characters, and can be composed of letters, numbers or underscores. The *snap_name* must be unique on the array.

-b

Identifies a Snapshot to be snapped.

-o

This option specifies an object name. The object name must be a valid object name that is one of the following object types:

- A device name
- A file system mount point
- A Veritas volume name
- A drive letter

snapcli attach

Attaches a mount point with one of its Snapshot instances.

DESCRIPTION

This operation is used to attach a snapshot named *snap_name* on *object_name*. It is also optionally used to decide whether to rescan the operating system's devices before attempting the attach operation. If an *object_name* is specified, you can optionally specify an *object_type*. Specifying an *object_type* without specifying an *object_name* is not permitted. You must specify a *snap_name* to identify the Snapshot.

This operation is performed on the host that has access to the mount point to be attached. You must perform an attach operation for each device you wish to access. If no *object_name* is given, all devices are searched and those belonging to the specified Snapshot are subsequently attached.

SYNTAX

```
snapcli attach -h | -s snap_name | [[-o object_name
[-t object_type]] | [-d driveletter]] [-r 0|1]
[-c rescan_count] [-f]
```

OPTIONS

-h

Prints out this message. If this option is present on the command line, no other option validation is performed and this message is printed.

-s

Specifies a snapshot name. This is a string of up to 256 characters, and can be composed of letters, numbers, or underscores. The *snap_name* must be unique on the array.

-o

Specifies an *object_name*. The *object_name* must be a valid object name of one of the following object types:

- A device
- A file system mount point
- A Veritas volume
- A drive letter

This option cannot be used in conjunction with the -d option.

-t

This option specifies the object type. This option is only a valid option when a -o option is given. The only valid object types are "device", "filesystem", "volume", and "driveletter". This option is optional for these four basic supported object types. If no option is provided, SnapCLI attempts to determine the correct object type among the four object types and reports a failure if it cannot.

This option cannot be used in conjunction with the -d option.

-d

This option is only valid for the Microsoft Windows platform and is ignored on all other platforms. This option specifies the drive letter to attempt to assign to the operating system device attached. If the attach operation performed results in more than one device attached, then only the first device in that list will be affected. If you wish to provide a drive letter for all devices attached, they must be attached individually.

This option cannot be used in conjunction with the -o or -t options.

- `-r`
 This option specifies whether to perform a rescan of the operating system's devices before performing an attach operation. The default rescanning behavior when `-r` is not present depends on whether an *object_name* was specified with the `-o` option. If an *object_name* is given, then the default behavior is to not perform a rescan operation. The default behavior is to perform a rescan of the operating system's devices.
 When present, `-r (0|1)` overrides default behavior.
- `-c rescan_count`
 This option specifies the number of rescans to be performed of the operating system's devices, after performing a successful attach operation to bring the devices online. If a count of 0 is specified, no rescans are performed. However, this could result in the device attaching successfully on the array, but not being online within the host.
 If the `-c` is not present, then a default value of 1 is set.
- `-f`
 This option forces SnapCLI to attach the snapshot, regardless of the current value of the snapshot's `allowReadWrite` property. A snapshot created with SnapCLI always has `allowReadWrite=No`, to prevent accidental modifications to the snapshot. The `-f` option does not modify the value of the `allowReadWrite` property, which can only be changed using Unisphere.

snapcli detach

Detaches a mount point from a snapshot.

DESCRIPTION

This operation is used to detach a snapshot named *snap_name* on *object_name*. If an *object_name* is specified, you can optionally specify an *object_type*. Specifying an *object_type* without specifying an *object_name* is not permitted. You must specify a *snap_name* to identify the snapshot.

This operation is performed on the host that has access to the mount point that will be detached. You must perform a detach operation for each device you wish to detach. If no *object_name* is provided, all devices will be searched and those belonging to the specified Snapshot will be subsequently detached.

SYNTAX

```
snapcli detach -h | -s snap_name [-o object_name [-t object_type]] [-f]
```

OPTIONS

- `-h`
 Prints out this message. If this option is present on the command line, no other option validation is performed and this message is printed.
- `-s`
 Specifies a snapshot name. This is a string of up to 256 characters, and can be composed of letters, numbers or underscores. The *snap_name* must be unique on the array.
- `-o`
 Specifies an *object_name*. The *object_name* must be a valid object name of one of the following object types:
- A device
 - A file system mount point

- A Veritas volume
- A drive letter

-t

This option specifies the object type. This option is only a valid option when a -o option is given. The only valid *object_type* are "device", "filesystem", "volume", and "driveletter". This option is optional for these four basic supported object types. If no option is provided, SnapCLI attempts to determine the correct object type among these four object types and reports a failure if it cannot.

-f

This option is only valid for the Microsoft Windows platform and this option is ignored on all other platforms. This option specifies whether to force SnapCLI to remove the drive letter after a detach operation.

CHAPTER 14

EMC Secure Remote Support on the Storage Processor

This chapter introduces EMC Secure Remote Support (ESRS) on the Storage Processor (SP) and the associated commands for configuring and managing this feature.

Major topics include:

- [ESRS Device Client on Storage Processor](#).....630
- [esrsconfig -agentprovision](#) 630
- [esrsconfig -agentstatus](#)..... 631
- [esrsconfig -configproxy -addproxy](#)..... 632
- [esrsconfig -configproxy -list](#).....632
- [esrsconfig -configproxy -removeproxy](#).....633
- [esrsconfig -configpolicymanager -addpm](#)..... 633
- [esrsconfig -configpolicymanager -list](#)..... 634
- [esrsconfig -configpolicymanager -addpolicyproxy](#)..... 634
- [esrsconfig -configpolicymanager -removepolicyproxy](#)..... 635
- [esrsconfig -configpolicymanager -removePM](#).....635

ESRS Device Client on Storage Processor

The ESRS device client on Storage Processor feature is included only in VNX operating environment (OE) for Block versions 5.32 that are later than version 05.32.000.5.209 or versions 05.33 that are later than version 05.33.000.5.051. This software monitors the operation of your VNX for Block systems for error events and automatically sends ConnectEMC notifications to your service provider. It also provides a path for your service provider to use to securely connect to your specified VNX for Block system (through the associated storage processor).

This solution offers a secure architecture from end to end, including the following features:

- EMC issues X.509 digital certificates to authenticate the ESRS device client on Storage Processor to EMC.
- EMC professionals are authenticated using two unique factors.
- All EMC service professionals have a unique username that is logged with all their actions.
- All communication originates from the Storage Processor. The ESRS device client on Storage Processor does not accept unsolicited connections from EMC or the Internet.
- All communications between EMC and the ESRS device client on Storage Processor includes the latest security practices and encryption technologies, including certificate libraries based on RSA Lockbox technology, and Advanced Encryption Standard (AES) 256-bit encryption.
- Those who implement the ESRS device client on Storage Processor solution can further control remote access by using the Policy Manager. The Policy Manager gives full control of how EMC interacts with VNX systems. SSL is available between the ESRS device client on Storage Processor and the Policy Manager.

For more information about the ESRS device client on Storage Processor feature for VNX for Block systems, refer to the *EMC Secure Remote Support for VNX* technical module on the EMC Online Support website at <http://Support.EMC.com>.

esrsconfig -agentprovision

Provisions ESRS on the storage processor (SP).

PREREQUISITES

The VNX system must be configured with at least one DNS server for EMC Secure Remote Support (ESRS) to function properly. To provision ESRS, you must be logged in with Administrator privileges.

DESCRIPTION

The `esrsconfig -agentprovision` command allows you to provision ESRS on the Storage Processor (SP).

SYNTAX

```
esrsconfig -agentprovision -user user [-emcuser] -password
password [-reprovision] [-http|-socks -proxyhost proxyhostname|
proxyipaddress -proxyport proxyportnumber | [-proxyuser
proxyhostuser -proxypassword proxypassword]]
```

OPTIONS

-user

- Specifies an EMC Online Support username. For EMC users, specify an EMC username.
- password
Specifies an EMC Online Support password. For EMC users, specify an RSA SecureID token password.
- emcuser
Specified when the user is an administrator or authorized service provider.
- reprovision
Specified if reprovisioning ESRS.
- http | -socks
Sets the proxy type.
- proxyhost
Specifies the proxy Host or IP address. Used only if either `-http` or `-socks` is specified, otherwise ignored.
- proxyport
Specifies the proxy port. Used only if either `-http` or `-socks` is specified, otherwise ignored.
- proxyuser
Specifies the proxy user name. Used only if either `-http` or `-socks` is specified, otherwise ignored.
- proxypassword
Specifies the proxy user password. Used only if either `-http` or `-socks` is specified, otherwise ignored.

EXAMPLE

```
naviseccli -h 10.244.70.103 -user administrator -password
passwordsample12 -scope 0 esrsconfig -agentprovision -user
brad.pitt@company.com -password Password1
```

OUTPUT

```
Software installed successfully.
```

```
ESRS has been provisioned and Template Call_Home_Template_7.33.0.9
applied successfully.
```

POSTREQUISITES

Provision ESRS for the other SP immediately after provisioning it for the first SP. This ensures ESRS will send CallHomes appropriately in the event of an SP failover.

esrsconfig -agentstatus

Returns the status of the ESRS device client.

DESCRIPTION

The `esrsconfig -agentstatus` command returns the status of the ESRS device client on the storage processor.

SYNTAX

```
esrsconfig -agentstatus
```

OUTPUT

```
Connecting To: esrs-corestg.emc.com on port 443
Connection Status: Connected
Proxy Server: Disabled
```

```
Policy Manager: Disabled.
SSL: Enabled, strength 168
Remote Sessions:
```

esrsconfig -configproxy -addproxy

Adds proxy server configuration information for ESRS on the SP.

DESCRIPTION

The `esrsconfig -configproxy -addproxy` command is used to add proxy server configuration information for ESRS on the SP.

SYNTAX

```
esrsconfig -configproxy -addproxy -http|-socks -proxyhost
proxyhostname | proxyipaddress -proxyport proxyportnumber |
[-proxyuser proxyhostuser -proxypassword proxypassword]
```

OPTIONS

```
-http | -socks
    Sets the proxy type.
-proxyhost
    Specifies the proxy Host or IP address. Used only if either -http or -socks is
    specified, otherwise ignored.
-proxyport
    Specifies the proxy port. Used only if either -http or -socks is specified,
    otherwise ignored.
-proxyuser
    Specifies the proxy username. Used only if either -http or -socks is specified,
    otherwise ignored.
-proxypassword
    Specifies the proxy user password. Used only if either -http or -socks is
    specified, otherwise ignored.
```

EXAMPLE

```
naviseccli -h 10.244.70.103 -user administrator -password
passwordexample1 -scope 0 esrsconfig -configproxy -addproxy -http
-proxyhost 10.245.36.219 -proxyport 80
```

OUTPUT

```
Agent Proxy configuration Successful.
```

esrsconfig -configproxy -list

Lists the proxy configuration status of ESRS on the SP.

DESCRIPTION

The `esrsconfig -configproxy -list` command lists the current status and details of the proxy server configuration for ESRS for the SP.

SYNTAX

```
esrsconfig -configproxy -list
```


OUTPUT

```
DRMProxy Server
Enabled: true
Proxy Type: HTTP
Host: 10.245.36.219
Port: 80
User name:
```

esrsconfig -configproxy -removeproxy

Removes proxy server configuration for ESRS for the SP.

DESCRIPTION

The `esrsconfig -configproxy -removeproxy` command removes the proxy server configuration for ESRS on the SP.

SYNTAX

```
esrsconfig -configproxy -removeproxy
```

OUTPUT

```
Agent Proxy configuration Successful.
```

esrsconfig -configpolicymanager -addpm

Adds and configures policy manager for ESRS on the SP.

DESCRIPTION

The `esrsconfig -configpolicymanager -addpm` command configures Policy Manager for EMC Secure Remote Support (ESRS) on the storage processor (SP).

SYNTAX

```
esrsconfig -configpolicymanager -addpm [-secure] -host
pmhostname | pmipaddress [-port PMportnumber] [-http|-socks
-proxyhost proxyhostname|proxyhostipaddress
[-proxyport proxyportnumber] [-proxyuser proxyhostuser
-proxypassword proxypassword]
```

OPTIONS

`-addpm`

Adds Policy Manager and configuration information.

`-secure`

If specified, uses the HTTPS protocol to access the server.

`-host`

Specifies the IP address of the policy manager host.

`-port`

Specifies the port used to access the policy manager.

`-http | -socks`

Sets the proxy type.

`-proxyhost`

Specifies the proxy Host or IP address. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxyport`

Specifies the proxy port. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxyuser`
Specifies the proxy username. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxypassword`
Specifies the proxy user password. Used only if either `-http` or `-socks` is specified, otherwise ignored.

EXAMPLE

```
naviseccli -h 10.241.216.50 -user sysadmin -password sysadmin -scope 0
esrsconfig -configpolicymanager -addpm -host 10.245.36.202
```

OUTPUT

```
Agent PM changes successful
```

esrsconfig -configpolicymanager -list

Lists the status of the policy manager configuration for ESRS on the SP.

DESCRIPTION

The `esrsconfig -configpolicymanager -list` command lists the details of the policy manager configuration for ESRS on the SP.

SYNTAX

```
esrsconfig -configpolicymanager -list
```

EXAMPLE

```
naviseccli -h 10.241.216.50 -user admin -password passwordsample1
-scope 0 esrsconfig -configpolicymanager -list
```

OUTPUT

```
Policy Server
  Enabled: true
  Host: 10.245.36.202
  Port: 8090
  SSL Enabled: false
Proxy Server
  Enabled: false
```

esrsconfig -configpolicymanager -addpolicyproxy

Adds proxy configuration to policy manager for ESRS on the SP.

DESCRIPTION

The `esrsconfig -configpolicymanager -addpolicyproxy` command adds the proxy server configuration information.

SYNTAX

```
esrsconfig -configpolicymanager -addpolicyproxy -http|-socks
-proxyhost proxyhostname | proxyhostipaddress [-proxyport
proxyportnumber] [-proxyuser proxyhostuser -proxypassword
proxypassword]
```

OPTIONS

```
-addpolicyproxy
```

Adds proxy server configuration to Policy Manager for ESRS on the SP.

`-http` | `-socks`
Sets the proxy type.

`-proxyhost`
Specifies the proxy Host or IP address. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxyport`
Specifies the proxy port. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxyuser`
Specifies the proxy username. Used only if either `-http` or `-socks` is specified, otherwise ignored.

`-proxypassword`
Specifies the proxy user password. Used only if either `-http` or `-socks` is specified, otherwise ignored.

EXAMPLE

```
naviseccli -h 10.241.216.50 -user admin -password passwordsample1
-scope 0 esrsconfig -configpolicymanager -addpolicyproxy
-socks -proxyhost 10.245.36.219
```

OUTPUT

```
Agent PM changes successful
```

esrsconfig -configpolicymanager -removepolicyproxy

Removes the proxy server for Policy Manager for ESRS on the SP.

DESCRIPTION

The `-configpolicymanager -removepolicyproxy` command removes the proxy server configuration from Policy Manager for ESRS on the SP.

SYNTAX

```
esrsconfig -configpolicymanager -removepolicyproxy
```

EXAMPLE

```
naviseccli -h 10.241.216.50 -user admin -password passwordsample2
-scope 0 esrsconfig -configpolicymanager -removepolicyproxy
```

OUTPUT

```
Agent PM changes successful
```

esrsconfig -configpolicymanager -removePM

Removes Policy Manager configuration from ESRS on the SP.

DESCRIPTION

The `esrsconfig -configpolicymanager -removePM` command removes the policy manager configuration from ESRS on the SP.

SYNTAX

```
esrsconfig -configpolicymanager -removePM
```

EXAMPLE

```
naviseccli -h 10.241.216.50 -user sysadmin -password sysadmin -scope 0  
esrsconfig -configpolicymanager -removePM
```

OUTPUT

```
Agent PM changes successful
```

CHAPTER 15

Data at Rest Encryption

This chapter introduces Data at Rest Encryption (D@RE) and the associated commands for activating and monitoring this feature.

Major topics include:

- [Data at Rest Encryption overview](#)..... 638
- [securedata -feature -activate](#)..... 639
- [securedata -feature -info](#)..... 640
- [securedata -backupkeys -status](#)..... 641
- [securedata -backupkeys -retrieve](#)..... 642
- [securedata -auditlog -retrieve](#)..... 643
- [securedata -auditlog -cksum](#)..... 644

Data at Rest Encryption overview

Data at Rest Encryption (D@RE) is provided through controller-based encryption (CBE) at a physical disk drive level. A unique data encryption key (DEK) is generated for each drive and is used to encrypt data as it is sent to the drive. The goal of this feature is to ensure that all customer data and identifying information will be encrypted with strong encryption, primarily to ensure security in the event of loss of a disk drive.

Note

Some unencrypted data could be in the system partition (for example, hostnames, IP addresses, dumps, and so on). In addition, there is potential for small amounts of unencrypted user data as a result of writing diagnostic materials to the system partition. All the data written to the array by using regular I/O protocols (iSCSI, FC) are encrypted. Anything that comes into the array by using the control path will not be encrypted by this solution; however, information that is sensitive (for example, passwords) are encrypted by a different mechanism (as they are on non-encrypting arrays).

For new VNX systems that are ordered with the D@RE feature, encryption should be enabled on the systems during manufacturing. Verify whether D@RE has been enabled and activated by running the `securedata -feature -info` command. If **Encryption Mode** appears as **N/A**, you need to perform a non-disruptive upgrade (NDU) of the DataAtRestEncryption enabler and activate it as described in the *Security Configuration Guide for VNX*. If **Encryption Mode** appears as **Unencrypted**, you only need to activate it using the `securedata -feature -activate` command.

NOTICE

Once activated, the encryption operation cannot be reverted. When possible, enable encryption prior to populating the system with data, RAID groups, and such. This action will avoid the data in place upgrade process and its effects on system cache and system performance.

For VNX systems that do not have D@RE enabled, enabling of encryption on the system requires a non-disruptive upgrade (NDU) of the DataAtRestEncryption enabler. This upgrade can be done upon request. A subsequent activate operation must be initiated through the `securedata -feature -activate` command.

A new component, referred to as the VNX Key Management Server, is responsible for generating, storing and otherwise managing the encryption keys for the system. The keystore that is generated to store the encryption keys resides on a managed LUN in private space on the system. Keys are generated or deleted in response to notifications that a RAID group/disk drive have been respectively added or removed.

Changes to the configuration of the system that result in changes to the keystore will generate alerts that recommend key backups be created. When an operation that results in a change to the keystore occurs, an alert will appear and persist until the keystore has been retrieved from the system for backup. Backup the keystore by using the `securedata -backupkeys -retrieve -path` command.

In the event that the keystore becomes corrupted, the system will be nonfunctional. The system will enter a degraded state, only the operating system boots. In this state, attempts to access the system through Unisphere will return an error indicating that the keystore is in an inaccessible state. In this case, a service engagement is required for resolution.

A separate auditing function is provided for general key operations that track all key establishment, deletion, backup, and restore changes as well as SLIC addition.

For additional information about the Data at Rest Encryption feature, refer to the *EMC VNX2: Data at Rest Encryption* white paper.

securedata -feature -activate

Initiates the encryption operation.

PREREQUISITES

In order to activate Data at Rest Encryption, the DataAtRestEncryption enabler must be installed on the system. A user role of administrator, storageadmin, or sanadmin is required to activate encryption. FAST cache must be destroyed before activating encryption.

DESCRIPTION

The `securedata -feature -activate` command activates the encryption operation.

SYNTAX

```
securedata -feature -activate [-status]
```

OPTIONS

`-status`

Returns the status of the current encryption operation. Note that the encryption activation request can take several minutes.

EXAMPLE #1

This example shows output for when the encryption enabler is not yet installed.

```
naviseccli -h <array_ip> securedata -feature -activate
```

```
Encryption operations are not supported until the
DataAtRestEncryption enabler is installed.
```

EXAMPLE #2

This example shows output for when the encryption feature is activated and no errors are returned.

```
naviseccli -h <array_ip> securedata -feature -activate
```

```
WARNING: You are about to activate the Data at Rest Encryption
Feature on the targeted array. This is not a revertible operation.
This will cause data encryption keys to be created and all user data
will begin to be encrypted. It is recommended that you have an up-to-
date and verified backup of your array as well as an up-to-date
config capture, created via the arrayconfig command, before executing
this command. Do you wish to continue (y or n)? y
```

EXAMPLE #3

This example shows output for when encryption cannot be activated because FAST cache is enabled on the array.

```
naviseccli -h <array_ip> securedata -feature -activate
```

```
Please make sure the FAST cache is disabled on your array. Please use
naviseccli command: cache -fast -destroy to disable FAST cache. After
you activate encryption you may recreate your FAST Cache using the
cache -fast -create command.
```

EXAMPLE #4

This example shows output for when the `-status` switch is used and encryption was successful.

```
naviseccli -h <array_ip> securedata -feature -activate -status
```

```
Status: Operation completed successfully
```

EXAMPLE #5

```
naviseccli -h <array_ip> securedata -feature -activate -status
```

This example shows output for when the `-status` switch is used and encryption has failed.

```
SP A: Activation failed because one or more BE I/O modules cannot support Data at Rest Encryption. Examine other SP Event Log entries for details, gather SPcollects and contact your service provider. (0x7140801b)
```

securedata -feature -info

Returns feature information for Data at Rest Encryption.

PREREQUISITES

This command is valid for all user roles except VASA Administrator.

DESCRIPTION

The `securedata -feature -info` command allows you to see the status and details for the Data at Rest Encryption feature.

SYNTAX

```
securedata -feature -info [-isencryptionsupported] [-encryptionmode] [-encryptionstatus] [-encryptionpct]
```

OPTIONS

`-isencryptionsupported`

Returns whether the encryption feature is available/enabled.

`-encryptionmode`

Returns the current encryption mode of the array.

`-encryptionstatus`

Returns the status of the encryption process. A status of `Scrubbing` indicates that the feature is in the process of wiping any clear text data that existed on unbound drives (drives previous removed from a RAID Group) when the encryption feature was activated.

`-encryptionpct`

Returns the percent of completion of the encryption process on the array.

EXAMPLE #1

This example shows output for when the encryption enabler is installed, but the encryption feature has not yet been activated.

```
naviseccli -h <array_ip> securedata -feature -info
```

```
Is Encryption Supported: Yes
Encryption Mode: Unencrypted
Encryption Status: Not encrypting
Percent Encrypted: 0.00
```


EXAMPLE #2

This example shows output for an encryption operation that is currently in progress.

```
naviseccli -h <array_ip> securedata -feature -info
```

```
Is Encryption Supported: Yes
Encryption Mode: Controller Based Encryption
Encryption Status: In progress
Percent Encrypted: 25.67
```

EXAMPLE #3

This example shows output for when an encryption operation is 100% complete, but a new drive has been added to the array, which must be scrubbed.

```
naviseccli -h <array_ip> securedata -feature -info
```

```
Is Encryption Supported: Yes
Encryption Mode: Controller Based Encryption
Encryption Status: Scrubbing
Percent Encrypted: 100.00
```

EXAMPLE #4

This example shows output for an encryption operation that is 100% complete.

```
naviseccli -h <array_ip> securedata -feature -info
```

```
Is Encryption Supported: Yes
Encryption Mode: Controller Based Encryption
Encryption Status: Encrypted
Percent Encrypted: 100.00
```

EXAMPLE #5

This example shows output for when the `-encryptionmode` switch is used. This will return only the encryption mode.

```
naviseccli -h <array_ip> securedata -feature -info -encryptionmode
```

```
Encryption Mode: Controller Based Encryption
```

securedata -backupkeys -status

Retrieves the status of the keystore backup operation.

PREREQUISITES

All user roles, with the exception of VASA administrator (vmadmin), can execute this command.

DESCRIPTION

The `securedata -backupkeys -status` command returns the status of the keystore backup operation. It also indicates whether the keystore has changed and if a backup operation is required.

SYNTAX

```
securedata -backupkeys -status
```

EXAMPLE #1

This example shows the output for when the encryption enabler is installed, but the encryption feature has not yet been activated on the array.

```
naviseccli -h <array_ip> securedata -backupkeys -status
```

```
Status: No operation required
```

EXAMPLE #2

This example shows the output for when a backup operation is required. This can display for the initial keystore backup operation, or because the keystore has changed since the last backup.

```
naviseccli -h <array_ip> securedata -backupkeys -status
```

```
Status: Backup keystore operation required
```

EXAMPLE #3

This example shows the output for when the backup operation was initiated and is in progress.

```
naviseccli -h <array_ip> securedata -backupkeys -status
```

```
Status: Backup keystore operation in progress
```

EXAMPLE #4

This example shows the output for when no backup operation is required because no additional changes to the keystore were made since the last backup completed.

```
naviseccli -h <array_ip> securedata -backupkeys -status
```

```
Status: Backup keystore operation complete
```

securedata -backupkeys -retrieve

Retrieves a backup file containing a copy of the keystore from the array.

PREREQUISITES

You must have a user role of administrator or securityadmin to execute this command.

DESCRIPTION

The `securedata -backupkeys -retrieve` command retrieves a backup file containing a copy of the keystore.

SYNTAX

```
securedata -backupkeys -retrieve -path backupfilepath
```

EXAMPLE #1

This example shows output for when the encryption enabler is installed, but the feature has not yet been activated.

```
naviseccli -h <array_ip> securedata -backupkeys -retrieve
```

```
This operation is not supported until the encryption feature has been activated.
```

EXAMPLE #2

This example shows output for when the file is retrieved successfully.

```
naviseccli -h <array_ip> securedata -backupkeys -retrieve -path C:
\Temp
```

```
The process of retrieving the file(s) is in progress
Successfully retrieved:
FNM00125200242_2014_03_10_15_21_14_0000000000000005.lbb
```

EXAMPLE #3

This example shows output for when the file retrieval fails because the request has timed out.

```
naviseccli -h <array_ip> securedata -backupkeys -retrieve -path C:
\Temp
```

```
The process of retrieving the file(s) is in progress
Operation failed
The retrieval process is taking longer than expected. The system is
too busy to handle the request. Please try again later.
```

securedata -auditlog -retrieve

Retrieves the audit log file and its accompanying checksum file from the array.

PREREQUISITES

You must have a user role of administrator or securityadmin to execute this command.

DESCRIPTION

The `securedata -auditlog -retrieve` command retrieves the audit log file and its accompanying checksum file from the array and stores them in a location specified by the user.

SYNTAX

```
securedata -auditlog -retrieve all | mmyyyy -path logpath
```

OPTIONS

all

Returns all log entries up to a limited file size of 100 MBs. If the log entries exceed 100 MBs, the first 100 MBs of the file will be returned (beginning with transaction ID 1). In these cases, the file name will contain the word "partial". Approximately 1,500,000 transactions are required to generate a file that exceeds 100 MBs. If a partial file is retrieved using the *all* option, utilize the *mmyyyy* option to retrieve the remaining log entries, starting at the month and year of the last log entry in the partial file.

mmyyyy

If specified, returns the log entries for a specific month and year.

-path

Specifies the path on the local system where the auditlog and checksum files will be saved.

EXAMPLE #1

This example shows output if the command succeeds, but with a partial file retrieval as the full log file exceeds 100 MBs.

```
naviseccli -h <array_ip> securedata -auditlog -retrieve all -path C:
\SecureData_Files
```

```
The process of retrieving the file(s) is in progress
Successfully retrieved:
```

```
FNM00125200242_2014_03_10_18_48_16_00000000000001B7_0000000000000789_p
artial.log
FNM00125200242_2014_03_10_18_48_16_00000000000001B7_0000000000000789_p
artial.hsh
```

EXAMPLE #2

This example shows input where the month and year were specified, and output if the command succeeds with a full log file retrieval for that month and year.

```
naviseccli -h <array_ip> securedata -auditlog -retrieve 032014 -path
C:\SecureData_Files
```

```
The process of retrieving the file(s) is in progress
Successfully retrieved:
FNM00125200242_2014_03_10_18_48_16_0000000000000001_00000000000002A_f
ull.log
FNM00125200242_2014_03_10_18_48_16_0000000000000001_00000000000002A_f
ull.hsh
```

EXAMPLE #3

This example shows output if the command fails because an invalid month/year combination is specified.

```
naviseccli -h <array_ip> securedata -auditlog -retrieve 031844 -path
C:\SecureData_Files
```

```
Invalid time period entered for the audit log retrieval request.
```

securedata -auditlog -cksum

Retrieves a regenerated checksum file from the array for a previously retrieved audit log file.

PREREQUISITES

You must have a role of administrator or securityadmin to execute this command.

A checksum file must have been previously generated using the `securedata -auditlog -retrieve` command.

Note

The newly retrieved checksum file will have a .hsh extension and the same file name as the original checksum file. To keep the original checksum file, you must save a copy of the original file with either a new name or to a different location so the original file will not be overwritten. If renaming a copy of the original file, it is recommended that you preserve the file-naming convention, which includes the timestamp and transaction ID range, and append the file name with a unique identifier, such as "_xxx".

DESCRIPTION

The `securedata -auditlog -cksum` command retrieves a regenerated checksum file for a previously retrieved audit log file.

SYNTAX

```
securedata -auditlog -cksum logFilename -path cksumPath
```

OPTIONS

```
-cksum
```

Specifies the file name of the audit log file for which to retrieve the regenerated checksum file.

`-path`

Specifies the path on the local system where the regenerated checksum file will be saved.

EXAMPLE #1

This example shows output if the command is executed successfully.

```
naviseccli -h <array_ip> securedata -auditlog -cksum
FNM00125200242_2014_03_10_18_48_16_0000000000000001_00000000000002A_f
ull.log -path C:\Temp
```

```
The process of retrieving the file(s) is in progress
Successfully retrieved:
FNM00125200242_2014_03_10_18_48_16_0000000000000001_00000000000002A_f
ull.hsh
```

EXAMPLE #2

This example shows output if the command fails. This can occur when the file requested to be retrieved is for a different array, or if the filename of the originally retrieved .log file was modified and the system cannot locate the requested log entries based on the filename.

```
naviseccli -h <array_ip> securedata -auditlog -cksum
CF2U6124800059_2014_05_21_18_02_33_0000000000000001_000000000000038_f
ull.log -path C:\SecureData_Files
```

```
The process of retrieving the file(s) is in progress
Operation failed
SP A: The checksum file retrieval process failed. Examine other SP
Event Log entries
for details, gather SPcollects and contact your service provider.
(0x71408019)
```

EXAMPLE #3

This example shows output if the command fails because the request has timed out.

```
naviseccli -h <array_ip> securedata -auditlog -cksum
FNM00125200242_2014_03_10_18_48_16_0000000000000001_00000000000002A_f
ull.log -path C:\SecureData_Files
```

```
The process of retrieving the file(s) is in progress
Operation failed
The retrieval process is taking longer than expected. The system is
too busy to
handle the request. Please try again later.
```

