



# **EMC<sup>®</sup> Navisphere<sup>®</sup> Command Line Interface (CLI)**

## **REFERENCE**

**For Basic, Access Logix<sup>™</sup>, MetaLUN, Reserved LUN Pool,  
iSCSI, LUN Migration, Domain and Security,  
Virtual Provisioning<sup>™</sup>, Event Monitor, and Virtual Server  
Commands**

**P/N 300-003-628  
REV A11**

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

Copyright © 2006–2009 EMC Corporation. All Rights Reserved.

Published August, 2009

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

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

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

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

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

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

<b>Preface .....</b>	<b>11</b>
<b>Chapter 1    <b>About EMC Navisphere CLI</b></b>	
Terminology .....	18
About EMC Navisphere software .....	22
About Navisphere Storage System Initialization Utility .....	22
About Navisphere Host and SP Agents .....	23
About Navisphere Server Utility .....	23
About Navisphere Manager .....	25
About Navisphere Storage Management Server .....	26
About Navisphere CLI .....	27
About Navisphere Analyzer .....	28
About Navisphere Quality of Server Manager .....	29
Navisphere Environments .....	29
Storage-system configuration and management with the CLI ..	31
Installing Navisphere CLI .....	31
Setting up the host or SP agent .....	31
Configuring a storage system with the CLI .....	31
Configuring and managing remote mirrors .....	32
AX4-5 series and AX series support .....	33
<b>Chapter 2    <b>About CLI Commands</b></b>	
About Secure CLI .....	38
Getting started with Secure CLI .....	39
Establishing a Secure CLI Security File .....	39
naviseccli .....	41
LUN IDs, unique IDs, and disk IDs .....	46

## Chapter 3 Basic Commands

Storage setup and management overview .....	52
Shared storage setup and management overview .....	54
alpa -get .....	56
alpa -set .....	57
arraycommpath .....	59
arrayconfig -capture .....	61
arrayconfig -duplicate .....	63
arrayname .....	66
backendbus -analyze .....	67
backendbus -get -speeds .....	71
backendbus -get -sfstate .....	73
backendbus -get -physical .....	74
backendbus -get -all .....	75
backendbus -resetspeed .....	77
baseuuid .....	79
bind .....	80
cachecard -initialize .....	90
cachecard -list .....	91
chglun .....	93
chrg .....	99
clearlog .....	101
clearstats .....	102
-cmdtime .....	103
copytohot spare .....	104
createrg .....	105
emconfiguration .....	108
failovermode .....	109
faults -list .....	111
firmware .....	113
flash -ioport .....	115
flash -iomodule .....	117
flashleds .....	118
getagent .....	119
getall .....	122
getarrayuid .....	156
getcache .....	158
getconfig .....	163
getcontrol .....	165
getcrus .....	168
getdisk .....	174
getlog .....	182
getloop .....	184
getlun .....	186

getresume .....	198
getrg .....	203
getsniffer .....	207
getsp .....	210
getsptime .....	212
getunusedluns .....	213
inserttestevent .....	214
ioportconfig -list .....	215
ioportconfig -persist .....	218
luncache -clear .....	219
luncache -list .....	221
managedby .....	222
managefiles -delete .....	223
managefiles -list .....	225
managefiles -retrieve .....	226
ndu -list .....	228
ndu -install .....	230
ndu -runrules .....	235
ndu -commit .....	237
ndu -status .....	238
ndu -revert .....	239
networkadmin -get .....	241
networkadmin -mib .....	246
networkadmin -route .....	248
networkadmin -set .....	251
port .....	255
powersaving .....	268
rebootSP .....	269
rebootpeerSP .....	270
remoteconfig -getconfig .....	271
remoteconfig -setconfig .....	273
removerg .....	275
resetandhold .....	276
responsetest .....	277
sc_off .....	281
setcache .....	283
setfeature .....	287
setsniffer .....	288
setspstime .....	290
setsptime .....	291
setstats .....	292
shutdown .....	294
shutdownsp .....	295

	shutdownpeersp .....	296
	spcollect -info .....	297
	spcollect -set.....	298
	spcollect.....	299
	spportspeed .....	301
	storagegroup.....	305
	systemtype .....	321
	trespass .....	323
	unbind .....	324
	unitserialnumber .....	325
<b>Chapter 4</b>	<b>LUN Expansion (MetaLUN) Commands</b>	
	MetaLUNs overview .....	328
	CLI metaLUN command functions.....	330
	metalun -destroy .....	331
	metalun -expand .....	332
	metalun -info .....	341
	metalun -list.....	343
	metalun -modify .....	350
	Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN .....	352
<b>Chapter 5</b>	<b>Reserved LUN Pool Commands</b>	
	Introduction to the reserved LUN pool.....	356
	Allocation of the reserved LUN pool .....	356
	Estimating the reserved LUN pool size.....	357
	Estimating a suitable reserved LUN pool size .....	357
	reserved -lunpool -addlun.....	359
	reserved -lunpool -list .....	361
	reserved -lunpool -rmlun .....	364
<b>Chapter 6</b>	<b>iSCSI Commands</b>	
	About iSCSI storage systems in IP networks.....	367
	Using iSNS within an iSCSI storage environment.....	370
	iSCSI commands overview.....	372
	connection -adduser .....	374
	connection -deleteuser .....	376
	connection -getuser .....	378
	connection -getport .....	380
	connection -setport .....	384
	connection -delport.....	386

connection -route.....	387
connection -pingnode .....	389
connection -tracertoute .....	391
connection -setshredauth .....	393
connection -delshredauth.....	395
connection -getshredauth.....	396
connection -addset .....	397
connection -delset.....	399
connection -modifyset .....	400
connection -getset.....	402
connection -addpath .....	404
connection -delpath .....	406
connection -modifypath .....	407
connection -verifypath.....	409
isns -addserver.....	411
isns -deleteserver.....	413
isns -listserver .....	415
isns -setprimary .....	416

## Chapter 7 LUN Migration Commands

migrate -start.....	420
migrate -cancel.....	423
migrate -modify.....	424
migrate -list .....	425

## Chapter 8 Domain and Security Commands

About domains and user accounts .....	429
About LDAP/Active Directory .....	431
Domain and security commands overview .....	432
domain -add.....	434
domain -list .....	436
domain -remove .....	437
domain -setmaster.....	439
ntp -list.....	441
ntp -set .....	443
security -adduser.....	446
security -changeuserinfo .....	448
security -list.....	450
security -listrole .....	451
security -rmuser.....	452
security -ldap -addserver .....	453

security -ldap -modifyserver.....	456
security -ldap -removeserver .....	459
security -ldap -listserver .....	460
security -ldap -addrolemapping.....	462
security -ldap -modifyrolemapping.....	464
security -ldap -removeolemapping .....	465
security -ldap -listrolemapping .....	466
security -ldap -synchronize .....	467

## Chapter 9 Virtual Provisioning Commands

Virtual Provisioning overview.....	470
storagepool -create.....	473
storagepool -list.....	475
storagepool -destroy.....	478
storagepool -modify .....	479
storagepool -expand.....	481
storagepool -cancelexpand.....	482
thinlun -create .....	483
thinlun -list .....	485
thinlun -modify .....	489
thinlun -destroy.....	490
thin -info.....	491

## Chapter 10 Event Monitor Commands

Event Monitoring overview .....	494
-portal -create .....	496
-portal -list.....	497
-portal -migrate .....	498
-portal -destroy.....	499
-template -list.....	500
-template -import.....	501
-template -export.....	502
-template -destroy.....	503
-template -findconflict.....	504
-template -swap.....	505
-monitor -addsystem .....	506
-monitor -removesystem.....	507
-monitor -listmapping.....	508
-monitor -applytemplate .....	509
-monitor -stoptemplate .....	510
-monitor -getlog .....	511



-monitor -enableresponse.....	512
-monitor -disableresponse .....	513
-monitor -getlogsize .....	514
-monitor -setlogsize .....	515
-monitor -insertttestevent.....	516
-monitor -responsetest.....	517
-monitor -reloadconfig .....	520

## Chapter 11 Virtual Server Commands

Virtual server overview .....	522
server -volmap .....	523
server -update .....	526
server -getagent .....	527
server -remoteconfig.....	528
server -register .....	531
hypervisor -addvirtualcenter .....	532
hypervisor -modifyvirtualcenter .....	533
hypervisor -removevirtualcenter.....	534
hypervisor -addesx .....	535
hypervisor -modifyesx .....	536
hypervisor -removeesx.....	537
hypervisor -listesx .....	538

## Appendix A CLI Error Codes

Secure CLI error codes.....	540
Storage group command error codes .....	551
SnapView, MirrorView, and other command error codes .....	553
Feature command error codes.....	566
General feature command errors.....	566
Mapped RAID errors.....	567
LUN migration errors.....	567
Connection and iSNS (iSCSI) errors .....	568
Domain and security errors.....	568
Thin provisioning errors .....	569
Event Monitor errors .....	574
Virtual Server CLI errors.....	575

## Appendix B Secure CLI Command Coverage

Command coverage .....	578
------------------------	-----

**Appendix C Navisphere Server Utility and Initialization Utility**

Using the server utility command line to generate a high-availability report .....	590
Preparing for software installation on a storage system ...	591
After uploading the HA report.....	593
Using the command line initialization utility.....	594
 <b>Index .....</b>	 <b>597</b>

---

*As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this manual may not be supported by all revisions 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 manual, please contact your EMC representative.*

### **Audience**

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

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

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

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

<b>Organization</b>	This manual contains eleven chapters and three appendixes as follows:
<a href="#">Chapter 1</a>	Introduces Navisphere and the CLI.
<a href="#">Chapter 2</a>	Introduces <b>naviseccli</b> command.
<a href="#">Chapter 3</a>	Explains the basic CLI commands, that is, those that are not specifically designed to manage optional features and storage group CLI commands.
<a href="#">Chapter 4</a>	Explains the LUN expansion (metaLUN) CLI commands.
<a href="#">Chapter 5</a>	Explains how to configure the reserved LUN pool and describes the CLI commands that you can use if SnapView software is installed on your system.
<a href="#">Chapter 6</a>	Explains the iSCSI commands.
<a href="#">Chapter 7</a>	Explains the LUN migration commands.
<a href="#">Chapter 8</a>	Explains the domain and security commands.
<a href="#">Chapter 9</a>	Explains the virtual provisioning commands.
<a href="#">Chapter 10</a>	Explains the event monitor commands.
<a href="#">Chapter 11</a>	Explains the virtual server commands.
<a href="#">Appendix A</a>	Lists the error codes you might receive from CLI commands.
<a href="#">Appendix B</a>	Lists the commands supported/not supported on various storage systems.
<a href="#">Appendix C</a>	Explains the CLI commands for Navisphere Server and Navisphere Initialization tools.

---

**Note:** CLI commands for the EMC MirrorView/Synchronous software are explained in the MirrorView/Synchronous CLI manual; CLI commands for the EMC MirrorView/Asynchronous software are explained in the MirrorView/Asynchronous CLI manual; CLI commands for the EMC SAN Copy software are explained in the SAN Copy CLI manual; and CLI commands for the EMC SnapView software (for both snapshots and LUN clones) are explained in the SnapView admsnap and CLI manual.

---

## Storage systems no longer covered in this document

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

Storage system removed	Last revision including the storage system
CX200, CX400, and CX600	300-003-628 Rev. A07
FC series; C series	069-001038 Rev. A14

## Related documentation

Related documents include:

For CX4 series, CX3 series, and CX series storage systems, refer to the latest version of the *EMC Navisphere Command Line Interface (CLI) Reference* (P/N 300-003-628), *EMC SAN Copy Command Line Interfaces Reference* (P/N 069001189), *EMC SnapView Command Line Interfaces (CLI) Reference* (P/N 069001181), *EMC MirrorView/Asynchronous Command Line Interface (CLI) Reference* (P/N 300-001-335), and *EMC MirrorView/Synchronous Command Line Interface (CLI) Reference* (P/N069001184).

For the most current management and security content for CX4 series, CX3 series, and CX series storage systems, refer to the EMC Navisphere Manager online help.

## Conventions used in this guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

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



### **CAUTION**

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

## Typographical conventions

This manual uses the following format conventions:

<b>This typeface</b>	Indicates text (including punctuation) that you type verbatim, all commands, pathnames, filenames, and directory names. It indicates the name of a dialog box, field in a dialog box, menu, menu option, or button.
<i>This typeface</i>	Represents variables for which you supply the values; for example, the name of a directory or file, your username or password, and explicit arguments to commands.
This typeface	Represents a system response (such as a message or prompt), a file or program listing.
[ ]	Encloses optional entries.
	Separates alternative parameter values; for example: <i>LUN-name</i>   <i>LUN-number</i> means you can use either the LUN-name or the LUN-number.

## Finding current information

The most up-to-date information about the EMC Navisphere CLI is posted on the EMC Powerlink® website. We recommend that you download the latest information before you run the CLI commands.

To access EMC Powerlink, use the following link:

<http://Powerlink.EMC.com>

After you log in, select **Support > Technical Documentation and Advisories** and find the following:

- ◆ Navisphere Host Agent/CLI and Utilities Release Notes
- ◆ The latest version of this manual that is applicable to your software revision
- ◆ *Installation Roadmap for CLARiiON Storage Systems*, which provides a checklist of the tasks that you must complete to install your storage system in a storage area network (SAN) or direct connect configuration.

## Where to get help

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

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

<http://Powerlink.EMC.com>

**Technical support** — For technical support, go to EMC WebSupport on Powerlink. To open a case on EMC WebSupport, you must be a WebSupport customer. Information about your site configuration and the circumstances under which the problem occurred is required.

### Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to [techpubcomments@EMC.com](mailto:techpubcomments@EMC.com) with your opinions of this guide.





This manual includes content for CX4 UltraFlex™ series (CX4-120 systems, CX4-240 systems, CX4-480 systems, and CX4-960 systems), CX3 UltraScale™ series (CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems, and CX3 model 80 systems), CX series, AX4-5 series, and AX series storage systems. For FC series and C series storage systems, refer to older versions of the *EMC Navisphere Command Line Interface (CLI) Reference* (P/N 069001038 Rev. A14) available in the Powerlink.

This chapter describes the EMC® Navisphere® storage-system management configurations and architecture.

Major topics are:

- ◆ Terminology ..... 18
- ◆ About EMC Navisphere software ..... 22
- ◆ Storage-system configuration and management with the CLI ... 31
- ◆ AX4-5 series and AX series support ..... 33

## Terminology

Term	Meaning
Access Logix™ software	Storage group option. See <b>shared storage system</b> .
Active Directory (AD)	Microsoft's centralized directory implementation.
Agent	EMC Navisphere® Agent that runs on a server; see also host agent.
AX series storage systems	AX150SC, AX100SC, AX150SCi, AX100SCi, AX150, AX100, AX150i and AX100i storage systems, where AX150SC, AX100SC, AX150SCi and AX100SCi are single-SP storage systems; AX150, AX100, AX150i and AX100i are dual-SP storage systems.
AX4-5 series storage systems	AX4-5, AX4-5i, AX4-5SC, AX4-5SCi storage systems, where AX4-5SC and AX4-5SCi are single-SP storage systems; AX4-5 and AX4-5i are dual-SP storage systems.
Back-end port	Ports associated with drive connections.
CDE	EMC CLARiiON® Driver Extensions failover software.
CHAP	Challenge Handshake Authentication Protocol is an authentication mechanism used by iSCSI targets to authenticate initiators or by iSCSI initiators to authenticate targets (mutual CHAP).
CLI	EMC Navisphere Command Line Interface.
CX4 series storage systems	CX4 model 120 systems, CX4 model 240 systems, CX4 model 480 systems, and CX4 model 960 systems.
CX3 series storage systems	CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems and CX3 model 80 storage systems.
CX series storage systems	CX200, CX300 series, CX400 series, CX500 series, CX600, CX700 storage systems.
domain	In the context of Navisphere software, a group of storage systems that you can manage from a single management application session. You can divide your storage systems into multiple domains as long as each storage system is in only one domain. Each domain has a directory that defines the storage systems in the domain.
EMC Navisphere Manager	The EMC Navisphere Manager application.
EMC Navisphere Express	The EMC Navisphere Express application.
ESX Server	The ESX server (hypervisor) is the hardware and software component of VMware that allows its users to create virtual machines in software.
Event Monitor	EMC Navisphere Event Monitor.

FLARE® software	Formerly known as Core or Base software.
Front-end port	Ports associated with external initiators (for example, servers, SAN copy™, MirrorView™ initiators).
HBA	Host bus adapter.
host agent	The EMC Navisphere Agent that runs on a storage-system server.
IQN	iSCSI Qualified Name is an iSCSI port worldwide name such as <code>iqn.1992-04.com.emc:cx.apm00034901792.a0</code> .
iSCSI	Internet SCSI protocol that uses the Internet Protocol (IP) for communication between iSCSI initiators and iSCSI targets.
iSCSI initiator	An iSCSI device that initiates requests to iSCSI targets such as a server that contains either NICs or iSCSI HBAs.
iSCSI target	An iSCSI device that responds to requests issued by initiators. For example, an iSCSI storage system is a target to iSCSI initiators.
managed agent	Host agent or SP agent managed by EMC Navisphere management software.
Lightweight Directory Access Protocol (LDAP)	A network protocol for accessing information in a directory.
LDAPS	Secure LDAP network protocol.
LDAP server	The external server that hosts the directory service that Navisphere will use to authenticate the users.
managed host	A host agent that you selected to manage.
managed storage system	A storage system managed by EMC Navisphere Manager.
Manager	See <a href="#">“EMC Navisphere Manager”</a>
NAS	Network attached storage.
Navisphere Express	A limited web-based management interface supported on AX4-5 series and AX series (also known as Navi Express).
persist	Writes configuration information to persistent memory. Sometimes this is referred to as initializing the port.
PowerPath	EMC PowerPath® failover software.
physical port ID	The location of the I/O port relative to its I/O module. The physical location is required to find the back-end ports on a CX4 series storage system.
physical slot ID	The slot that is inserted into the I/O module. The physical location is required to find the back-end ports on a CX4 series storage system.

RAID group storage system	A storage system whose SPs are running FLARE software with RAID group functionality.
role mapping	A process for assigning the Navisphere access level to an external server user group or user. The Navisphere access levels are Monitor, Manager, Administrator, and Security Administrator.
SAN	Storage area network.
secret	A password used for CHAP authentication in an iSCSI environment.
server or managed server	A host with a managed storage system.
shared storage system	A storage system with the EMC Storage Group (Access Logix) option, which provides data access control (storage groups) and configuration access control. A shared storage system is always a RAID group storage system.
single-SP system	A storage system that has one SP. See <a href="#">“AX4-5 series storage systems”</a> and <a href="#">“AX series storage systems”</a> .
SP	Storage processor.
SP agent	The Navisphere Agent that runs on an SP.
storage group	A collection of one or more LUNs that you select, and to which you can connect one or more servers.
storage pool	General term used for RAID groups and thin pools. In the Navisphere UI Storage tree, a storage pool is displayed as a single node beneath the storage system icon. The RAID Groups and Thin Pools icons are displayed beneath the Storage Pools icon.
Storage Management Server software	Software that runs in a storage system or portal server and lets a person manage it over a network.
thin LUN	A logical unit of storage that may have a subscribed user capacity that is greater than the user capacity of the shared thin pool. More than one thin LUN can share the user capacity of the same thin pool. The storage capacity of a thin LUN grows by using the storage in the shared thin pool A thin LUN name can consist of 1-64 characters. A thin LUN's minimum user capacity is 1GB, the maximum is 14 terabytes.
thin pool	A set of disks (maximum depends on storage system type), all with the same redundancy (RAID 5 and RAID 6 only), that shares its user capacity with one or more thin LUNs. For optimum efficiency, we recommend that all disks have the same capacity. RAID 6 is the default RAID type. You can expand the user capacity of a thin pool by adding disks to the thin pool.
user capacity	The total amount of physical storage capacity in the thin pool available to thin LUNs (raw capacity minus overhead).
RAID overhead	Thin pool space reserved for internal operations (not available to thin LUNs).

raw capacity	Total amount of physical storage capacity in the thin pool (user capacity plus overhead).
consumed capacity	The amount of user capacity consumed by the thin LUNs in the pool.
available capacity	User capacity minus consumed capacity.
percent full threshold	The consumed capacity of a thin pool that triggers an alert. The default value is 70%. We recommend that you set the value somewhere between 50 and 75%. Regardless of the value that you set, the storage system automatically triggers an alert when the consumed capacity reaches the 85%, 90% and 95% full threshold value and then with each 1% increase in consumed capacity (96%, 97%, 98%, 99% and 100%).
subscribed user capacity	The total amount of thin LUN user capacity that is configured in the thin pool and is presented to the attached hosts. The subscribed capacity can exceed the physical user capacity of the thin pool.
oversubscribed user capacity	The amount of subscribed user capacity that exceeds the user capacity of the thin pool.
percent subscribed	The percentage of thin pool user capacity that has been assigned to all thin LUNs in the pool and is visible to the hosts.
unshared storage system	A storage system without the EMC Storage Group option.
Virtual Provisioning™	A feature that provides thin provisioning, which allows you to assign more storage capacity to a host than is physically available. Thin provisioning uses thin pools and thin LUNs to allocate storage. With thin provisioning, storage is assigned to the server in a capacity-on-demand method from a shared storage pool (thin pool). The storage system software monitors and adds storage capacity, as required, to each thin pool, not each LUN. This simplifies the creation and allocation of storage capacity.
Virtual Center	Virtual Center is the management software that manages one or more ESX Servers.
Virtual Machine	The Virtual Machine (VM) is a virtual host created by software using a hypervisor. It is recognized as a standard physical host on a network. A Virtual Machine represents a complete system, with processors, memory, networking, storage and BIOS.
Windows	A supported version of Windows: Windows Server 2003, Windows XP, Windows Server 2008, and Vista.

---

## About EMC Navisphere software

The EMC Navisphere software suite is composed of the following components:

- ◆ Navisphere Storage System Initialization Utility
- ◆ Navisphere Host and SP Agents
- ◆ Navisphere Server Utility
- ◆ Navisphere Manager
- ◆ Navisphere Storage Management Server
- ◆ Navisphere CLI
- ◆ Navisphere Analyzer
- ◆ Navisphere Quality of Server Manager

---

### About Navisphere Storage System Initialization Utility

You must initialize the storage system in order to register the server HBAs and NICs with the storage system and have access to the storage system's data.

For CX4 series, CX3 series, and CX series Fibre Channel storage systems, use the utility to discover storage systems, and set network parameters (IP address, subnet mask, and default gateway). In addition, for CX3 series storage systems with iSCSI data ports, use the utility to set network parameters for these ports.

---

**Note:** For CX series storage systems, an authorized service provider must install and run the initialization utility.

---

For AX4-5 series and AX series storage systems, use the utility to discover storage systems, set network parameters (IP address, subnet mask, and default gateway), and create management user accounts. In addition, for AX150 iSCSI storage systems use the utility to set network parameters for the storage system's iSCSI data ports.

---

**Note:** The Navisphere Initialization Utility can be run on servers with connects to the management ports of CX4 series, CX3 series, CX700, CX500 series, CX300 series, AX4-5 series, and AX series storage systems. CX4 series systems must be running FLARE OE 04.28.xx.yy.z or later. CX3 series systems must be running FLARE OE 03.22.xxx.5.yyy or later. CX700, CX500 series, CX300 series, and AX100 series systems must be running FLARE OE 02.16.xxx.5.yyy or later. AX150 series systems must be running FLARE 02.20.xxx.5.yyy or later. AX4-5 series systems must be running FLARE 02.23.xxx.5.yyy.

---

---

## About Navisphere Host and SP Agents

Agents relay configuration and status requests from manager to CLARiiON® storage systems. Agents also periodically retrieve information from CLARiiON storage systems, and forward it on to Manager for update of the information it displays. Agents also monitor storage-system events and can notify personnel by email, page, or modem when any designated event occurs.

Host agents reside on attached hosts. SP agents reside on the storage processors (SPs), and are installed at the factory.

---

**Note:** We recommend that you always have at least one privileged user in the SP agent configuration file.

---

---

## About Navisphere Server Utility

Similar to the host agent, the Navisphere Server Utility registers the server's host bus adapter (HBA) or network interface card (NIC) with the attached CX4 series, CX3 series, CX series, AX4-5 series, or AX series storage system; however, unlike the host agent, the server utility does not:

- ◆ automatically retrieve information from storage systems. With the server utility, you must manually update the information by starting the utility or you can create a script to run the utility.

---

**Note:** For Windows servers running version 6.20 or later, you can automatically update server information if the **Registration Service** feature remains enabled after installing the server utility.

---

- ◆ send LUN mapping information back to a CX4 series, CX3 series or CX series storage system. This information will not display in the manager UI or CLI. To see LUN mapping information in the manager UI or CLI for a CX4 series, CX3 series or CX series storage system, you must use the host agent instead of the server utility.

---

**Note:** LUN mapping information is sent to AX4-5 series and AX series storage systems. In Navisphere Manager, this information displays in the manager UI or with the CLI **-lunmapinfo** command. In Navisphere Express, this information displays in the **Manage Virtual Disks** page.

---

- ◆ require network connectivity to the storage system.

---

**Note:** If you are using the server utility to upload a high-availability report to the storage system, you must have network connectivity.

---

- ◆ require that you install the utility; you can run it from the CD. However, we recommend that you install it on the server. Some features are not available unless you install the server utility on the server.

With the server utility, you can perform the functions listed below:

- ◆ Fibre Channel connections can determine if the server is configured for high availability (HA) by verifying that the server has at least one connection path to each storage-system SP, and that PowerPath® or some other failover software is running.
- ◆ Windows servers with Microsoft iSCSI initiators can use the utility to configure iSCSI connections - log on, log off, and remove an iSCSI target - and to configure mutual CHAP.
- ◆ AX4-5 series storage systems running Navisphere Express can use the utility to manage snapshots. If you are managing your AX series storage system with Navisphere Manager or if you have a CX series storage system, you must use SnapView™ and the admsnap utility to manage snapshots. For more information on SnapView and admsnap, refer to the *EMC SnapView Command Line Interfaces Reference*.

If you have the host agent installed but want to use some of the features of the server utility, you must install revision 6.22.20 or later of the server utility. The registration feature of the server utility will be disabled and the host agent will be used to register the server's HBAs to the storage system. Prior to revision 6.22.20 of the server



utility, you could not install the host agent and server utility on the same server. However, you could install them on different servers that are connected to the same storage system.

---

**Note:** The server utility runs on servers attached to CX4 series, CX3 series, CX series, AX4-5 series, and AX series storage systems. CX3 series systems must be running FLARE OE 03.22.xxx.5.yyy or later. CX series and AX100 series systems must be running FLARE 02.16.xxx.5.yyy or later. AX150 series systems must be running FLARE 02.20.xxx.5.yyy or later. AX4-5 series systems must be running FLARE 02.23.xxx.5.yyy. Not all features are supported in all revisions of the server utility. Refer to the server utility's online help for information on which features are available.

---

---

## About Navisphere Manager

Manager is a centralized storage-system management tool for configuring and managing CLARiiON storage systems. It provides the following basic functionality:

- ◆ Discovery of CLARiiON storage systems
- ◆ Status and configuration information display
- ◆ Event management
- ◆ Storage configuration and allocation

Manager is a web-based user interface that lets you securely manage CLARiiON storage systems locally on the same LAN or remotely over the Internet, using a common browser. Manager resides on a storage system or a supported version of Windows that is running the storage management server software, and is downloaded to the browser when the storage management server software is accessed.

---

**Note:** For an AX series storage system, you must complete an upgrade from Navisphere Express to Navisphere Manager. Refer to the document for upgrading a storage system to Navisphere Manager and the EMC Navisphere Manager online help.

---

---

## About Navisphere Storage Management Server

The Navisphere Storage Management Server software is provided with Navisphere Manager, and is installed and executed on each SP in a CX4 series, CX3 series, CX series, AX4-5 series, or AX series storage system, or optionally on a Windows XP, Windows server. A storage system with this software installed is a storage management server. A server with this software installed is called a server portal. All CX4 series, CX3 series, CX series (excluding the CX200 and CX300), AX4-5 series, and AX series storage systems leave the factory with the storage management server software installed on the SPs.

The storage management server software performs the following functions:

- ◆ Receives and responds to requests from Navisphere Manager
- ◆ Forwards requests to the local SP agents for processing
- ◆ Forwards status and configuration updates to Navisphere Manager
- ◆ Replicates user and domain information to all storage systems in domain
- ◆ Authenticates user logins and authorizes user requests
- ◆ Logs all user logins and actions

---

## About Navisphere CLI

The CLI complements and can be used as an alternative to Navisphere Manager. It provides a command line interface for storage-system management, including storage provisioning, status and configuration information retrieval, and control. You can use the CLI to automate management functions through shell scripts and batch files. CLI commands for many functions are server-based and are provided with the host agent. The remaining CLI commands are web-based and are provided with software that runs in the storage-system SPs. Navisphere CLI includes the Secure on supported operating systems. See [Chapter 2, "About CLI Commands."](#)

---

**Note:** The Navisphere CLI is supported on all CX4 series, CX3 series, CX series, AX4-5 series, and some AX series storage systems.

You must install the Navisphere Manager enabler on AX150 storage systems in order for Navisphere CLI to be operational. Refer to the document for upgrading a storage system to Navisphere Manager. Once you install the Navisphere Manager enabler, Navisphere Express is no longer operational. You cannot uninstall the Navisphere Manager enabler once you install it. You cannot use Navisphere Express and Navisphere Manager to co-manage an AX series storage system.

---

AX4-5 series storage systems support only Secure CLI.

The architecture of the Navisphere version 6.X Manager and CLI differ as follows:

- ◆ Navisphere Manager is a web-based product that communicates with the storage management server software that runs on SPs or Windows servers. The SPs and servers are grouped in entities called storage domains. When you log in to a storage system or server within the domain, Navisphere Manager provides a domain-based view in which you can manage your storage systems.
- ◆ Navisphere CLI 6.X is a command line interface that has some commands that are server-based and some that are web-based. Server-based commands communicate with the host agent. Web-based commands communicate with storage management server software. In the CLI, you issue commands to individual storage systems through a command line structure.

The tradeoffs between manager and the CLI are as follows:

Manager	CLI
Uses a user interface (UI) with online, context-sensitive help	Uses a command line interface with a single help file through which you find the topic you want
Requires user interaction	Uses command lines that you can type interactively or write into a shell script for automated operations

With the CLI, as with manager, you can configure, control, and retrieve status from any managed storage system on the LAN. You can also use the CLI to automate disk-storage management functions by writing shell scripts or batch files.

---

**Note:** If you are running VMware ESX Server, you can run Navisphere CLI from ESX Server or a virtual machine (VM). If you run Navisphere CLI from ESX Server, all commands are available. If you run Navisphere CLI from a VM, commands that require the Navisphere Host Agent are not supported, since the Navisphere Host Agent is not supported on a VM. The following commands require the Navisphere Host Agent: lunmapinfo, register. For commands that you can issue to either an SP or host agent, such as getagent or remoteconfig, if you issue the command to a VM, only the SP agent type is supported.

---

## About Navisphere Analyzer

Navisphere Analyzer 6.X is a web-based tool, using a common browser, that allows an administrator to graphically examine the performance characteristics of the logical and physical entities that make up a CLARiiON storage system. Analyzer supports immediate (real-time) data display, as well as the display of previously logged data. As a result, Navisphere Analyzer lets you do immediate comparisons, long-term trend analysis, and off-site performance troubleshooting/analysis.

Like Navisphere Manager, Navisphere Analyzer resides on a storage system or a supported version of Windows server that is running the storage management server software, and is downloaded to the browser when the storage management server software is accessed.

Analyzer is an optional product and requires an enabler to be installed in order to retrieve usable data.

---

## About Navisphere Quality of Server Manager

Navisphere Quality of Service Manager (NQM or Navisphere QoS Manager) lets you allocate storage-system performance resources on an application-by-application basis. You can use NQM to solve performance conflicts in consolidated environments where multiple applications share the same storage system. Within storage-system capacity, NQM lets you meet specific performance targets for applications, and create performance thresholds to prevent applications from monopolizing storage-system performance.

When you use NQM, you classify application I/O requests into user-defined I/O classes such as all I/O requests issued by an application or a particular host, in order to monitor performance and control the performance of this class of I/O. You specify various I/O characteristics of your applications to define your I/O classes, and set performance goals in order to meet your application's service level requirements.

NQM lets you centrally manage storage-system resources with no host footprint (no host-management or software installs required on the host).

---

## Navisphere Environments

You start Navisphere Manager in a supported browser on a Windows or UNIX host by entering the IP address of an SP in the storage system with Navisphere Manager installed on it. The Navisphere Manager user interface (UI) is downloaded to the Windows or UNIX host and it runs within the browser. All requests from the UI go to the storage management server software on the storage systems, which communicate with the SP agent. The SP agent, in turn, communicates with the FLARE or Storage Group software on the SPs.

Figure 1 shows a sample Navisphere Manager environment with Fibre Channel storage systems and an iSCSI storage system in a storage domain with various servers.

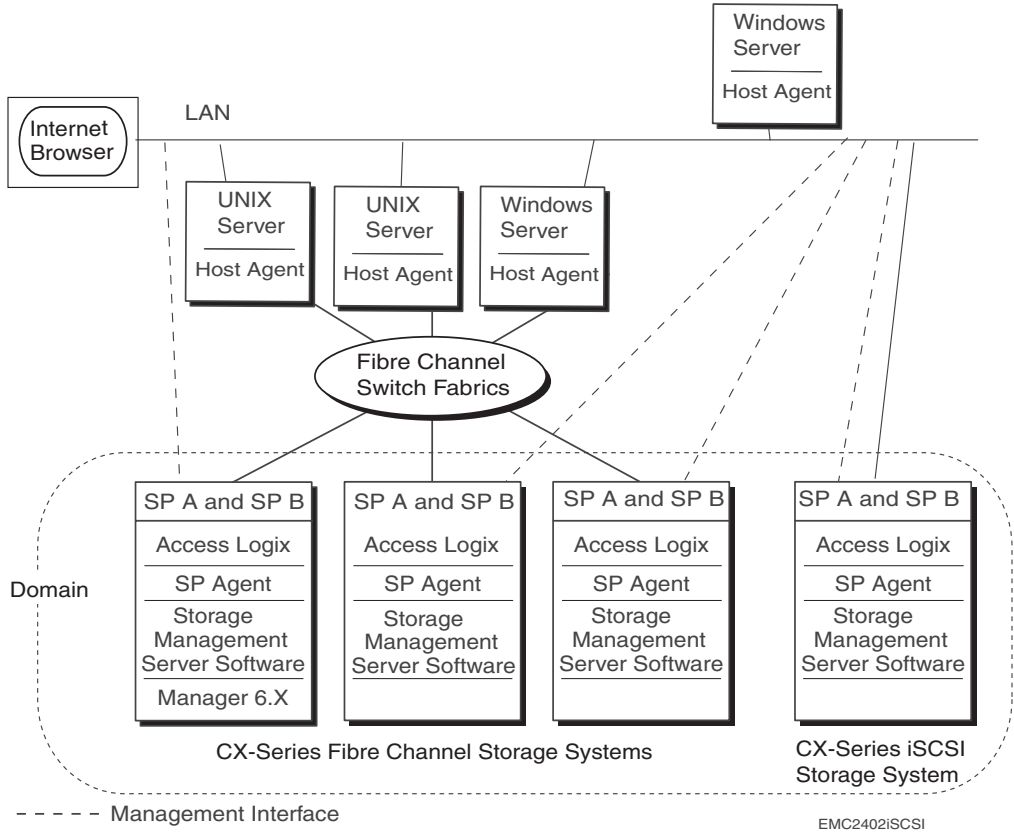


Figure 1 Fibre Channel and iSCSI storage systems in an IP network

---

## Storage-system configuration and management with the CLI

Before you can configure or manage shared storage systems with the CLI, you need to set up the Navisphere environment. For convenience with certain operations, you may want to use Navisphere Manager. Manager has a different architecture from the CLI, as explained on [page 27](#).



### **IMPORTANT**

Until you enable data access control for a storage system, any server connected to it can write to any LUN on it. To ensure that servers do not write to LUNs that do not belong to them, the procedures below assume that *either just one server is physically connected to the shared storage system or that just one server has been powered up since the servers were connected to the storage system*. You will use this server (called the configuration server) to configure the storage system.

---

---

### Installing Navisphere CLI

Install the CLI on the host you want to manage, as explained in the EMC CLARiiON server support products installation guide for your operating system.

---

### Setting up the host or SP agent

Configure the SP agent on each SP in the storage system, explained in the Navisphere Manager online help.

Configure the host agent on the server, explained in the EMC CLARiiON server support products installation guide for your operating system.

---

### Configuring a storage system with the CLI

For an unshared storage overview, see the section “[Storage setup and management overview](#)” on [page 52](#).

For a shared storage overview, see the section “[Shared storage setup and management overview](#)” on [page 54](#).

---

## Configuring and managing remote mirrors

Since the CLI communicates with the SP agent through one SP at a time, using the CLI to configure and manage a remote mirror is inefficient. We recommend that you use Navisphere Manager to do this because it can monitor both SPs on the primary and secondary storage systems and can also provide a more complete picture of the remote mirror configuration and status.



## AX4-5 series and AX series support

You can use Navisphere CLI with AX4-5 series and AX series storage systems. The CLI provides a command line interface for which you can write shell scripts to automate operations.

You must upgrade your storage system to Navisphere Manager before you can use Navisphere CLI with an AX series (AX150 series and AX100 series) storage system. Refer to the document for upgrading a storage system to Navisphere Manager.

---

**Note:** AX4-5 series storage systems running either Navisphere Express or Navisphere Manager can be managed with Secure CLI. Some commands are not available with Express.

You cannot use Navisphere Express and Navisphere Manager to co-manage an AX4-5 series and AX series storage system. Once you install Navisphere Manager on an AX4-5 series or AX series storage systems, Navisphere Express is no longer operational. You cannot uninstall Navisphere Manager once you install it.

---

We recommend that you always have at least one privileged user in the SP agent configuration file.

In general, commands supported by CX4 series, CX3 series, and CX series systems are supported by AX4-5 series and AX series systems. Because of some functionality differences, there are some commands that are unique to AX4-5 series and AX series systems and others that are not available. See the following table.

Command	Supported on		
	AX4-5 series with Navisphere Manager	AX4-5 series with Navisphere Express	AX series with Navisphere Manager
alpa	Yes <sup>a</sup>	No	Yes <sup>a</sup>
arraycomppath	Yes	No	No
arrayconfig	No	No	No
arrayname	Yes	Yes	Yes <sup>a</sup>
backendbus	No	No	No
baseuuid	No	No	No
bind	Yes <sup>a</sup>	Yes <sup>a</sup>	Yes <sup>a</sup>
cachecard	No	No	Yes
chglun	Yes <sup>a</sup>	Yes <sup>a</sup>	Yes <sup>a</sup>
chgrg	Yes <sup>a</sup>	Yes <sup>a</sup>	Yes <sup>a</sup>
clearstats	No	No	No
clearlog	Yes	Yes	Yes
failovermode	Yes	No	No
faults	Yes	Yes	Yes
firmware	No	No	No
getcrus	Yes <sup>a</sup>	Yes <sup>a</sup>	Yes <sup>b</sup>
getsniffer	No	No	No
getresume	Yes	Yes	Yes
getall	Yes	Yes	Yes
getlun	Yes	Yes	Yes
initializearray	No	No	No
managedby	Yes	Yes	Yes
metalun	Yes <sup>a</sup>	Yes	Yes <sup>a</sup>

Command	Supported on		
	AX4-5 series with Navisphere Manager	AX4-5 series with Navisphere Express	AX series with Navisphere Manager
migrate	Yes	Yes <sup>c</sup>	No
mirrorview	Yes	No	No
ndu	Yes	Yes <sup>a</sup>	Yes <sup>a</sup>
remoteconfig	No	No	Yes <sup>a</sup>
responsetest	Yes <sup>a</sup>	Yes <sup>a</sup>	Yes <sup>a</sup>
rebootpeersp	Yes	Yes	Yes <sup>d</sup>
rebootsp	Yes	Yes	Yes
sc_off	No	No	No
setcache	Yes	No	No
setsniffer	No	No	No
setstats	No	No	No
setspstime	Yes	Yes	Yes
sportspeed	Yes	Yes	Yes
shutdown	Yes	Yes	Yes
shutdownsp	Yes	Yes	No
shutdownpeersp	Yes	Yes	No
storagegroup	Yes	Yes	Yes <sup>a</sup>
systemtype	Yes	No	No
sportspeed	Yes	Yes	Yes

- a. Some of the switches used with this command are not supported.
- b. Some switches are supported only on AX series.
- c. Some switches are supported only on AX4-5 series.
- d. This command is supported only on AX150 series.

**Note:** Only Secure CLI commands are supported with AX4-5 series storage systems.

**What next?** Continue to the chapter listed below for the type of commands you want:

Basic (unshared storage), storage group (shared storage) [Chapter 3](#)

LUN expansion (metaLUN) [Chapter 4](#)

Reserved LUN pool [Chapter 5](#)

Internet SCSI (iSCSI) [Chapter 6](#)

LUN migration [Chapter 7](#)

Domain and security [Chapter 8](#)

Virtual Provisioning commands [Chapter 9](#)

Event Monitor commands [Chapter 10](#)

Virtual Server commands [Chapter 11](#)

---

This chapter explains the EMC Navisphere CLI commands.

**Note:** If you currently use Classic or Java CLI to issue CLI commands, please note that Secure CLI is replacing both Classic and Java CLI. Refer to the *Navisphere Command Line Interface (CLI) Reference* (P/N 300-003-628 Revision A02) for Java CLI support and *Navisphere Command Line Interface (CLI) Reference* (P/N 300-003-628 Revision A10) for Classic CLI support.

Major topics are:

- ◆ About Secure CLI ..... 38
- ◆ Getting started with Secure CLI ..... 39
- ◆ navisecli ..... 41
- ◆ LUN IDs, unique IDs, and disk IDs ..... 46

## About Secure CLI

Secure CLI is a comprehensive Navisphere CLI solution that provides one application and one security model for all CLI commands. Secure CLI provides role-based authentication, audit trails of CLI events, and SSL-based data encryption. Secure CLI is a thin client with both forward and backwards compatibility. The Secure CLI client does not need to be upgraded to work with later versions.

---

**Note:** Refer to the Host Agent/CLI and Utilities Release Notes, available on the Powerlink® website, for a list of supported operating systems. You must be running FLARE Operating Environment version 02.19.xxx.5.yyy or higher to use Secure CLI.

---

Secure CLI commands run in a command window. Each command consists of the **naviseccli** command (and switches) together with another subcommand (and its switches).

---

**Note:** Secure CLI provides a single summary of warnings and a single confirmation. Secure CLI preserves original command syntax and output for script compatibility.

---

Secure CLI does support commands equivalent to Classic CLI commands targeted directly at host agents, such as **-lunmapinfo** (see [“About Navisphere Host and SP Agents” on page 23](#) for information on host and SP agents). For commands that you can issue to either an SP or host agent, such as **-getagent**, if you issue the command to a host agent, Secure CLI displays an error message. You either must use Classic CLI or the appropriate Secure CLI command to issue commands to host agents.

---

**Note:** Secure CLI does not distinguish case of characters, so, regardless of the host operating system, you can use either uppercase, lowercase, or any combination of characters as you type commands.

---

If a Secure CLI command fails and the CLI does not generate its own error message, it displays an error message from the SP agent. Secure CLI generates errors about command line syntax for commands and options and their values. Secure CLI commands return 0 if the command is successful, or greater than zero if the command is unsuccessful.

---

## Getting started with Secure CLI

Before you begin to issue Secure CLI commands, you must create a Navisphere user account on the storage system. To create the required user accounts using Navisphere CLI, see [“Domain and Security Commands” on page 427](#). For details on using Navisphere 6.X security, refer to the *EMC Navisphere Manager help*.

You can also choose to configure a Navisphere 6.X security file to issue Secure CLI commands (see the next section, [“Establishing a Secure CLI Security File”](#)) on the host. If you establish a security file, you do not need to include the switches **-user**, **-scope**, and **-password** (or the password prompt), in each command you issue.

---

**Note:** Secure CLI requires the switches **-user**, **-scope**, and **-password** (or the password prompt) in each command line; you do not need to provide these switches in the command line if you establish a security file.

---

---

## Establishing a Secure CLI Security File

A storage system will not accept a command from Secure CLI unless the user who issues the command has a valid user account on the storage system. You can specify a valid account username, password, and scope (global or local) for each command you issue, or, more conveniently, you can create a Navisphere Security File.

The Navisphere security file is an encrypted file stored for each user on each host. You can add or delete a user security file using the **-AddUserSecurity** or **-RemoveUserSecurity** functions as arguments to the **naviseccli** command. You cannot copy a security file to another host. You must issue the **-AddUserSecurity** function on the host for which you want to create the security file.

When you create a security file, the username you use to log in to the current host is automatically stored in the security file, or you can specify an alternative username for the security file in the **-AddUserSecurity** request using the optional **-user** switch. If you omit the **-user** switch, the security file uses your current username. You can also target Security File at a specific storage-system using the **-ip** option.

For example, to add yourself to the security file on the current host, given the alternative username **altusername**, the password **mypass** and the scope **0** (global scope), type:

```
naviseccli -AddUserSecurity -password mypass -scope 0 -user altusername
```

Then, on this host, you can enter CLI commands to any storage system on which you have an account that matches the username **altusername**, with password **mypass** and global scope (**scope 0**).

You can also construct a security file with a default credential and security credentials for a specific storage-system. Multiple security credentials can be added to a Security File by specifying either a SP hostname or IP address when adding the credential. But only one credential can exist for a particular SP. The addition of a new default credential will cause the old one to be overwritten.

---

**Note:** Username and password are case sensitive.

---

The security file is stored in your default home directory. With Secure CLI, you can specify an alternative file path using the optional **-secfilepath** switch.

---

**Note:** If you specify an alternative location for the security file, you must specify the file path in every subsequent CLI command you issue to ensure the CLI locates the security file.

---

To save the example used above to the alternative location "c:\altlocation\" type:

```
naviseccli -AddUserSecurity -password mypass -scope 0 -user altusername - secfilepath c:\altlocation\
```

Then, for each subsequent command you issue, you must specify the **-secfilepath** switch with the security file path location "c:\altlocation\" in the command line.



## naviseccli

Sends status or configuration requests to a storage system from the command line

**Description** The **naviseccli** command sends storage-system management and configuration requests to a storage system from the Internet.

**Syntax** The **naviseccli** command is used as follows:

**naviseccli -help**

or

**naviseccli**

**[-address** *IPAddress* | *NetworkName* | **-h** *IPAddress* | *NetworkName*]

**[-AddUserSecurity]**

**[-f** *filename*]

**[-m]**

**[-nopoll** | **-np]**

**[-parse** | **-p]**

**[-password** *password*]

**[-port** *port*]

**[-q]**

**[-RemoveUserSecurity]**

**[-scope** 0 | 1 | 2]

**[-timeout** | **-t** *timeout*]

**[-user** *username*]

**[-v]**

**[-xml]**

*CMD* [*optional\_command\_switches*]

**Prerequisites** Anyone who can log in to a host running Navisphere CLI 6.X or later can use this command.

**Options** The **naviseccli** switches are:

**-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 follow instead.

**-address** *IPAddress | NetworkName* | **-h** *IPAddress | NetworkName*

Specifies the IP address or network name of the targeted SP on the desired storage system. The default, if you omit this switch, is **localhost**. This CLI version supports IPv6 network communication to the management ports of storage systems with FLARE version 04.28.xx.yy.z or later. It does not support communication with the data ports or communication between storage systems in a domain. Currently, you must still assign an IPv4 address to each SP, even for IPv6 configurations. For detailed information on IPv6, refer to the *Internet Protocol Version 6 (IPv6) Addressing Architecture, RFC 3513, April 2003*.

**-AddUserSecurity**

Directs the CLI to add user security information to the security file on this host. 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 (see **-password**), 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 host login name in the security file you create on this host. If you use the **-secfilepath** switch, you can specify an alternative location to your default home directory, for the security file on this host. You must then use the **-secfilepath** switch in each subsequent command you issue.

---

**-f** *filename*

Specifies to store data in a file.

**-m**

Suppresses output except for values. This option is most useful when used as part of a script.

---

**Note:** Only supported for commands that originated in Classic CLI.

---

**-nopoll** | **-np**

Directs the feature provider not to issue a poll request. This switch significantly increases performance when dealing with large or multiple storage systems. The feature provider automatically polls 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 storage 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 storage system. The default is 443. If you choose to change the default port number, management port 2163 will be supported; however, you must 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.

---

**Note:** This switch is supported only for commands that originated in Classic CLI.

---

### **-RemoveUserSecurity**

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

**-scope 0 | 1 | 2**

Specifies whether the user account on the storage system you want to log in to is local or global. A **0** (default) indicates global; a **1** indicates local; **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 storage systems in the domain.

A local account is effective on only the storage systems for which the administrator creates the account. The user can log in to only those storage systems on which he or she has a local account.

LDAP maps the username/password entries to an external LDAP or active directory server for authentication.

Username/password pairs whose roles are not mapped to the external directory will be denied access.

For authentication within the local security directory, specify global or local scope.

**-secfilepath *filepath***

Stores the security file in a file path location you specify. When you create a security file on a host 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.

---

**-timeout | -t *timeout***

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

**-user *username***

Specifies the username on the storage system you want to log in to. Usernames are case sensitive. You can omit this if your username has been added 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.

---

**Note:** Only supported for commands that originated in Classic CLI.

---

**-xml**

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

*CMD*

One of a set of commands used with the **naviseccli** command.

**CMD switches**

The *CMD* switches are described on the pages that follow.

**Spaces in arguments**

Normally, each argument to a CLI command consists of numbers, letters, and other valid printable characters for the operating system on which the CLI is running.

If a filename or other argument includes a space, you must enclose the entire string that includes the name in quotation marks ("xx xx").

## LUN IDs, unique IDs, and disk IDs

In Navisphere, the term *LUN ID* or *LUN number* means the unique integer assigned to the LUN when it is bound. When you bind a LUN, you can select the ID number. If you do not specify one, the default for the first LUN bound is 0, the second 1, the third 2, and so on. The maximum number of LUNs supported (and thus the valid range of LUN IDs) depends on storage system and operating system. See the EMC Support Matrix on the Powerlink website for the number of LUNs supported with your configuration.

The term *unique ID* applies to storage systems, SPs, HBAs, and switch ports. It means the World Wide Name (WWN) or World Wide Port Name (WWPN), which is a number designed to be unique in the world.

Storage systems have 16-byte unique IDs; SPs, HBAs, LUNs, and switch ports have 32-byte unique IDs. In storage systems that use storage groups, the Storage Group (Access Logix™) software assigns the unique IDs. Generally, unique IDs are transparent to users, but sometimes you need to specify them in commands. You can discover the unique IDs by entering the appropriate CLI **list** or **get list** command for the operation you want to perform.

*Disk IDs* depend on the kind of storage system.

CX3 series and CX series storage systems can have multiple back-end buses, and thus need a bus designator before the enclosure and disk designator:

*bus-number enclosure-number disk-number (b e d for short)*

For example, disk ID 000 (or 0\_0\_0) indicates the first bus or loop, first enclosure, and first disk, and disk ID 100 (1\_0\_0) indicates the second bus or loop, first enclosure, and first disk. A sample CX700 model storage system is shown in [Figure 2 on page 47](#).

---

**Note:** AX4-5 series storage systems have only one back-end bus (*bus-number=0*). AX series storage systems have only one back-end bus (*bus-number=0*) and only one enclosure (*enclosure-number=0*).

---

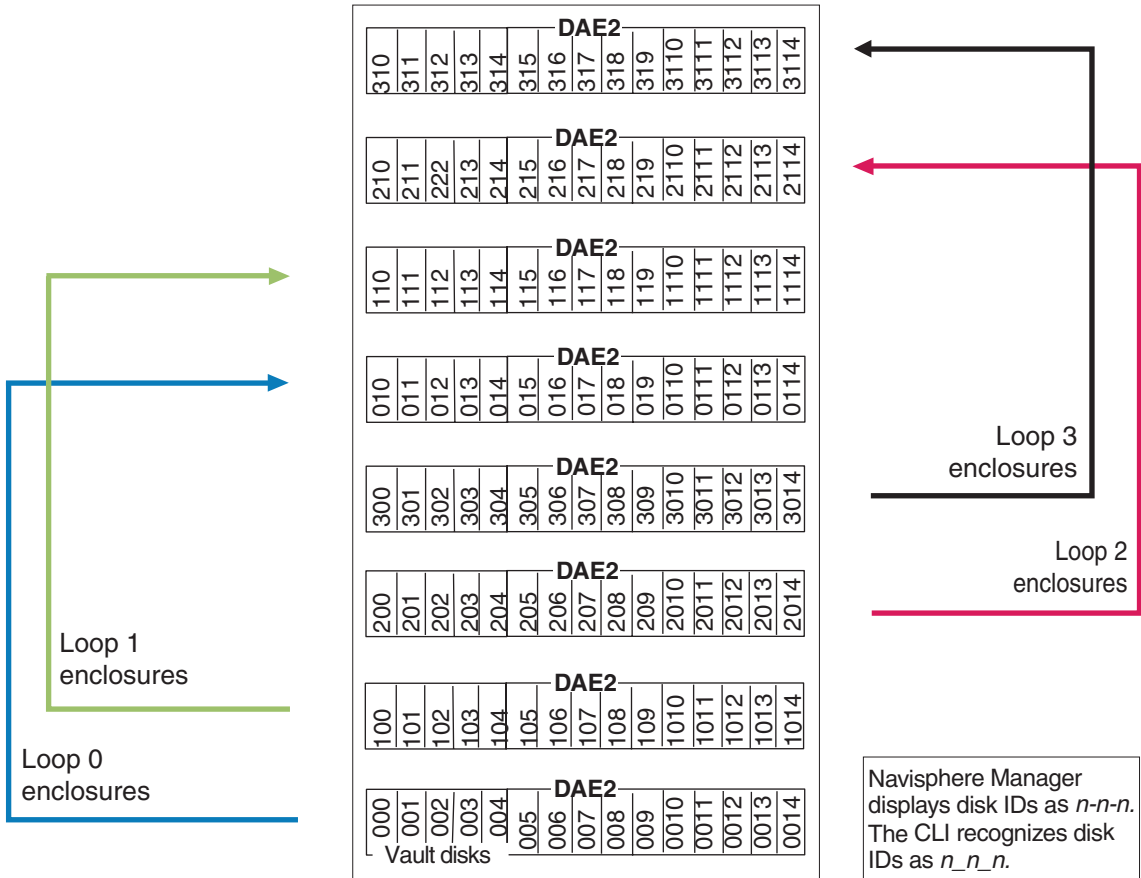


Figure 2 CX700 disk IDs

**What next?** Continue to the chapter listed below for the type of commands you want:

Basic (unshared storage) and storage group (shared storage) [Chapter 3](#)

LUN expansion (metaLUN) [Chapter 4](#)

Reserved LUN pool [Chapter 5](#)

Internet SCSI (iSCSI) [Chapter 6](#)

LUN migration [Chapter 7](#)

Domain and security [Chapter 8](#)

Virtual Provisioning (thin provisioning) [Chapter 9](#)

Event Monitor commands [Chapter 10](#)

Virtual Server commands [Chapter 11](#)



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

**Note:** CLI commands for the EMC MirrorView/Synchronous software are explained in the MirrorView/Synchronous CLI manual; CLI commands for the EMC MirrorView/Asynchronous software are explained in the MirrorView/Asynchronous CLI manual; CLI commands for the EMC SAN Copy software are explained in the SAN Copy CLI manual; and CLI commands for EMC SnapView (for both snapshots and LUN clones) are explained in the SnapView admsnap and CLI manual.

Major topics are:

◆ Storage setup and management overview .....	52
◆ Shared storage setup and management overview .....	54
◆ alpa -get .....	56
◆ alpa -set.....	57
◆ arraycommpath .....	59
◆ arrayconfig -capture .....	61
◆ arrayconfig -duplicate .....	63
◆ arrayname .....	66
◆ backendbus -analyze .....	67
◆ backendbus -get -speeds .....	71
◆ backendbus -get -sfpstate.....	73
◆ backendbus -get -physical .....	74
◆ backendbus -get -all.....	75
◆ backendbus -resetspeed .....	77

◆ baseuuid .....	79
◆ bind .....	80
◆ cachecard -initialize .....	90
◆ cachecard -list .....	91
◆ chglun .....	93
◆ chgrg .....	99
◆ clearlog .....	101
◆ clearstats .....	102
◆ -cmdtime .....	103
◆ copytohot spare .....	104
◆ createrg .....	105
◆ emconfiguration .....	108
◆ failovermode .....	109
◆ faults -list .....	111
◆ firmware .....	113
◆ flash -ioport .....	115
◆ flash -iomodule .....	117
◆ flashleds .....	118
◆ getagent .....	119
◆ getall .....	122
◆ getarrayuid .....	156
◆ getcache .....	158
◆ getconfig .....	163
◆ getcontrol .....	165
◆ getcrus .....	168
◆ getdisk .....	174
◆ getlog .....	182
◆ getloop .....	184
◆ getlun .....	186
◆ getresume .....	198
◆ getrg .....	203
◆ getsniffer .....	207
◆ getsp .....	210
◆ getsptime .....	212
◆ getunusedluns .....	213
◆ inserttestevent .....	214
◆ ioportconfig -list .....	215
◆ ioportconfig -persist .....	218
◆ luncache -clear .....	219
◆ luncache -list .....	221
◆ managedby .....	222
◆ managedby .....	222
◆ managefiles -delete .....	223

◆ managefiles -list .....	225
◆ managefiles -retrieve .....	226
◆ ndu -list .....	228
◆ ndu -install .....	230
◆ ndu -runrules .....	235
◆ ndu -commit .....	237
◆ ndu -status .....	238
◆ ndu -revert .....	239
◆ networkadmin -get .....	241
◆ networkadmin -mib .....	246
◆ networkadmin -route .....	248
◆ networkadmin -set .....	251
◆ port .....	255
◆ powersaving .....	268
◆ rebootSP .....	269
◆ rebootpeerSP .....	270
◆ remoteconfig -getconfig .....	271
◆ remoteconfig -setconfig .....	273
◆ removevg .....	275
◆ resetandhold .....	276
◆ responsetest .....	277
◆ sc_off .....	281
◆ setcache .....	283
◆ setfeature .....	287
◆ setsniffer .....	288
◆ setsptime .....	290
◆ setsptime .....	291
◆ setstats .....	292
◆ shutdown .....	294
◆ shutdownsp .....	295
◆ shutdownpeersp .....	296
◆ spcollect -info .....	297
◆ spcollect -set .....	298
◆ spcollect .....	299
◆ spsportspeed .....	301
◆ storagegroup .....	305
◆ storagegroup .....	305
◆ trespass .....	323
◆ unbind .....	324
◆ unitserialnumber .....	325

## Storage setup and management overview

This section shows how you might configure and monitor a storage system using **naviseccli** command. It shows commands with UNIX device names. The Windows device names would be drive letters. The storage system is a system, in which each SP is a host. For an overview showing storage systems with storage group, see [Section “Shared storage setup and management overview” on page 54](#). For storage systems without storage group:

1. Create four RAID groups.

```
naviseccli -h ss1_spa createrg 0 0_0_0 0_0_1
0_0_2 0_0_3 0_0_4
```

Creates RAID group 0 from disks 0 through 4 in the DPE enclosure. A new RAID group has no RAID type until it is bound.

```
naviseccli -h ss1_spa createrg 1 0_0_5 0_0_6
```

Creates RAID group 1 from disks 5 and 6 in the DPE enclosure.

```
naviseccli -h ss1_spa createrg 2 0_0_7 0_0_8
```

Creates RAID group 2 from disks 7 and 8 in the DPE enclosure.

```
naviseccli -h ss1_spa createrg 3 0_0_9
```

Creates RAID group 3 from disk 9 in the DPE enclosure.

---

**Note:** To create RAID groups for storage systems running Navisphere Express, you must use the **-raidtype** switch with the **createrg** command.

---

2. Bind a LUN on each RAID group.

```
naviseccli -h ss1_spa bind r5 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 ss1_spa bind r1 1 -rg 1
```

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 ss1_spa bind r1 2 -rg 2
```

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.

```
naviseccli -h ss1_spa bind hs -rg 3
```

Binds a hot spare on RAID group 2. The hot spare has no LUN ID (it is not really a LUN) and occupies all space on RAID group 2.

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

```
naviseccli -h ss1_spa getsp -mem
```

```
Memory Size For The SP:          930
```

Each SP has the same amount of memory, so you need not issue this command for SP B.

4. Set up storage-system caching with 70% memory for write caching and 30% for read caching.

```
naviseccli -h ss1_spa setcache -wcza 650 -wc 1 -p 8
-rcza 280 -rczb 280 -rc 1
```

This command sets the write cache for SP A (applies to both SPs) to 650 MB, enables write caching, sets a write cache page size of 8 KB, sets a read cache size for each SP of 280 MB, and enables read caching.

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

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

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

This command changes the name of the storage system to **users11**.

At this point, you can create file systems (run the appropriate host command) 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/or **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 Navisphere Manager than the CLI to create and manipulate storage groups. However, the following is a sequence of tasks you can perform using only the CLI:

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 52” on page 49](#) 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 manager or the CLI command **storagegroup -create**.
4. Connect the hosts to storage groups using manager 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 manager 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.

---

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

## alpa -get

### alpa command background

#### Displays the SCSI ID associated with an SP port

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 storage system, for ports 0 and 1, you can specify SCSI IDs 0 and 1, respectively. On the second storage system, for the ports you can specify IDs 2 and 3 respectively, and so on.

### Description

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** is used with **naviseccli** (described on [page 41](#)) as follows:  
**alpa -get**

### Prerequisites

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

### Output

See above. If the version of FLARE 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
```

### Example

```
naviseccli -h payroll_storage_spa alpa -get
Storage Processor:  SP A
Port ID:           0
ALPA Value:       0
```

For SP A (SP hostname **payroll\_storage\_spa**), this command lists the SCSI ID associated with port 0.



## alpa -set

Updates the SCSI ID associated with an SP port

### 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 storage 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` is used with `naviseccli` (described on [page 41](#)) as follows:

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

### Prerequisites

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

### Options

`-sp a | b`

Specifies the SP: A or B.

`-portid portid SCSI-ID`

Specifies the new SCSI ID to associate with the port. Port IDs range from 0 through 3 (CX3 series and CX series). Valid port SCSI IDs range from 0 through 126.

`-o`

Executes the command without prompting for confirmation.

**Output** See above. If the version of FLARE 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
```

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

```
This operation will cause a storage system reboot!  
Do you wish to continue(y/n)? y
```

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

## arraycommpath

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

### Description

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

This command works only for storage systems without storage group. For storage systems with storage group, use the command `storagegroup sethost -arraycommpath` (see [page 305](#)).

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 storage system inaccessible.**

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

### Syntax

`arraycommpath` is used with `naviseccli` (described on [page 41](#)) as follows:

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

### Prerequisites

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

### Options

**0**

Disables a communication path.

**1**

Enables a communication path.

**-o**

Executes the command without prompting for confirmation.

**Output** If the version of FLARE 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

**Example** `naviseccli -h ss1_spa arraycommpath 0`

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

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

## arrayconfig -capture

### Queries the storage system for configuration information

**Description** The **arrayconfig -capture** command queries the storage system for its configuration along with I/O port configuration information. When issued, the command will capture a storage 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 storage systems or rebuild the same storage system if the previous configuration is destroyed. The information captured may include:

- ◆ Software installed in the storage system (name of the software, revision, and so on).
- ◆ Physical configuration of the storage system (storage processors, disks, enclosures, and so on.)
- ◆ Logical configuration of the storage system (RAID groups, LUNs, storage groups, clones, SnapView, and mirrors).
- ◆ I/O port configuration information (hardware and logical configuration information) of a CX4 series storage system. You can find the hardware information in **CLARiiON > Physical > StorageProcessors > StorageProcessor > IODevices** and the logical configuration information in **CLARiiON > Logicals > IOPorts**.

**Syntax** **arrayconfig -capture** is used with **navisecli** (described on [page 41](#)) as follows:

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

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

**Options** **-format XML | CSV**

If **-format** is not specified, it defaults to XML.

**-logfile**

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

---

**Note:** All **-capture** commands can be run with monitor and above privileges.

---

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

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

## arrayconfig -duplicate

**Creates the components in the indicated configuration file on the targeted storage system**

### 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 storage 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 storage system qualifies for duplication:

#### Physical components

- ◆ The model number on the targeted storage system must be exactly the same as the source.
- ◆ The targeted storage 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 storage system must have all the system and Navisphere 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 storage system(s): Snapview sessions, SAN Copy information, synchronous/asynchronous mirror secondary information. If the XML captured from a mirror secondary storage system is used for duplication, then no mirror information will be created in the target storage 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 storage system; use it only after completely understanding the command's behavior.**

---

**Syntax** `arrayconfig -duplicate` is used with `naviseccli` (described on [page 41](#)) as follows:

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

**Prerequisites** Manager (or above) privileges are required to execute the `arrayconfig -duplicate` command.

**Options** Optional switches are:

**-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 logfile stored on the targeted storage system.

**-o**

Executes the command without prompting for confirmation.



**-recover**

Use this switch in a disaster recovery situation (if the storage 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 storage 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 have been 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 creates another RAID group with the same ID, the rollback will destroy the newly created raidgroup 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 and above privileges. All other **-duplicate** commands require manager or above privileges.

---

**Example** `naviseccli -h ss1_spa arrayconfig -duplicate -input c:\Capture158.xml`

## arrayname

### Renames a storage system

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

**Note:** AX4-5 series and AX series storage-system names cannot exceed 32 characters in length.

The CLI prompts for confirmation as follows.

**Syntax** `arrayname` is used with `naviseccli` (described on [page 41](#)) as follows:  
`arrayname` *NewArrayName*

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

**Options** `arrayname` *NewArrayName*

Specifies the new name for the storage system.

**Output** If the version of FLARE software running on the SP does not support this command, a `Not Supported` error message is printed to **stderr**.

**Example** `naviseccli -h ss1_spa arrayname users11`

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

This command changes the name of the storage system to **users11**.

## backendbus -analyze

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

**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 storage 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 storage system meets conditions required for a reset (see the list of conditions later in this section).

The current speed indicates the speed a bus or device is currently running at; the maximum speed indicates the highest speed a bus or device is capable of running at. 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.

Summary analysis <sup>a</sup>	Description
Speed <i>may</i> change on rescan	A speed-related fault exists on the back-end bus (see <b>Description</b> section).The CLI cannot determine whether a reset will occur.
Speed <i>will</i> change on rescan	A reset will occur.
Speed <i>will not</i> change on rescan	There are limiting devices on the back-end bus. A reset will <i>not</i> 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 appear 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 show the list of conditions.

Condition	Description (pass criteria)
Analyze Speed Rule	The summary analysis for the back-end bus, displayed in <b>backendbus -analyze</b> , indicates that a reset <i>will</i> or <i>may</i> 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** is used with **naviseccli** (described on [page 41](#)) as follows.

**backendbus -analyze**

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

## Output

Bus 0 - Current Speed: 2Gbps. Maximum speed: 4Gbps.

Device	Current Speed	Max Speed	Limiting	Type
-				
Bus 0 Enclosure 0	2	4	No	DAE3P
LCC A	2	4	No	
LCC B	2	4	No	
Disk 0	2	2	Yes	Fibre Channel
Disk 1	2	4	No	Fibre Channel
Disk 2	2	4	No	Fibre Channel
Disk 3	2	4	No	Fibre Channel
Disk 4	2	4	No	Fibre Channel
Disk 5	2	4	No	Fibre Channel
Disk 6	2	4	No	Fibre Channel
Disk 7	2	4	No	Fibre Channel
Disk 8	2	4	No	Fibre Channel
Disk 9	2	4	No	Fibre Channel
Disk 10	2	4	No	Fibre Channel
Disk 11	2	4	No	Fibre Channel
Disk 12	2	4	No	Fibre Channel
Disk 13	2	4	No	Fibre Channel
Disk 14	2	4	No	Fibre Channel

Bus 1 - Current Speed: 2Gbps. Maximum speed: 4Gbps.

Device	Current Speed	Max Speed	Limiting	Type
-				
Bus 1 Enclosure 1	2	4	No	DAE3P
LCC A	2	4	No	
LCC B	2	4	No	
Disk 0	2	2	Yes	Fibre Channel
Disk 1	2	4	No	Fibre Channel
Disk 2	2	4	No	Fibre Channel
Disk 3	2	4	No	Fibre Channel
Disk 4	2	4	No	Fibre Channel
Disk 5	2	4	No	Fibre Channel
Disk 6	2	4	No	Fibre Channel
Disk 7	2	4	No	Fibre Channel
Disk 8	2	4	No	Fibre Channel
Disk 9	2	4	No	Fibre Channel
Disk 10	2	4	No	Fibre Channel
Disk 11	2	4	No	Fibre Channel
Disk 12	2	4	No	Fibre Channel
Disk 13	2	4	No	Fibre Channel
Disk 14	2	4	No	Fibre Channel

```
Bus 0 is currently running at 2Gbps.  
Limiting Devices:  
Bus 0 Enclosure 0 Disk 0  
Speed will not change on rescan
```

```
Bus 1 is currently running at 2Gbps.  
Limiting Devices:  
Bus 1 Enclosure 1 Disk 0  
Speed will not change on rescan
```

```
Analyze Speed Rule...Failed No loop will change speed on rescan.  
System Faults Rule... Passed  
Active Sessions Rule... Passed  
Transitions Rule... Passed  
I/O Rate Rule... Passed
```

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

```
naviseccli -h ss1_spa backendbus -analyze
```

## backendbus -get -speeds

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

**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 storage 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 storage 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** is used with **naviseccli** (described on [page 41](#)) as follows.

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

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

**Options** *backend bus number*

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

**Output** Bus 2:

```
Current Speed: 2Gbps.  
Available Speeds:  
2Gbps.  
4Gbps.
```

**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
```



## backendbus -get -sfpstate

Displays the current condition of each back-end SFP

**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 storage 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` is used with `naviseccli` (described on [page 41](#)) as follows.

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

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

**Options** *backend bus number*

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

**Output** Bus 0:

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

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

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

## backendbus -get -physical

Displays physical slot and physical port ID information.

**Description** The **naviseccli backendbus** command with the **-get -physical** switches displays the physical location of the back-end bus on the CX4 series storage 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** is used with **naviseccli** (described on [page 41](#)) as follows:

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

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

**Options** *backend bus number*  
Specifies the back-end bus for which you want to display bus information.

**Example** **naviseccli -h ss1\_spa backendbus -get -physical**

## backendbus -get -all

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

**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 storage system. It displays the physical location of the back-end bus (I/O module and the physical port ID) only for the CX4 series storage 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 storage 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**.

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** is used with **naviseccli** (described on [page 41](#)) as follows.

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

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

**Options** *backend bus number*

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

**Output**

```
Bus 0

Current Speed: 4Gbps.
Available Speeds:
2Gbps.
4Gbps.
SPA SFP State: Online
SPB SFP State: Online
I/O Module Slot: 0
Physical Port ID: 0
```

```
Bus 1

Current Speed: 4Gbps.
Available Speeds:
2Gbps.
4Gbps.
SPA SFP State: Online
SPB SFP State: Online
I/O Module Slot: 1
Physical Port ID: 3
```

**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
```

## backendbus -resetspeed

Resets the back-end bus speeds for a storage system

**Description** The **naviseccli backendbus** command with the **-resetspeed** function, resets the back-end bus speed for the buses on the storage system. When you execute **backendbus -resetspeed**, if the back-end bus is eligible for a reset and the storage 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 storage 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 storage 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 <b>backendbus -analyze</b> to verify the following: <ul style="list-style-type: none"> <li>• back-end bus and device speed information</li> <li>• no limiting devices exist on the back-end bus</li> <li>• storage system meets required conditions</li> </ul>
2	Save the output from <b>backendbus -analyze</b> or <b>backendbus -get -speeds</b> , to compare the bus speed in the saved output to the bus speed after you initiate <b>backendbus -resetspeed</b> ,

**Syntax** **backendbus -resetspeed** is used with **naviseccli** (described on [page 41](#)) as follows.

**backendbus -resetspeed**

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

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

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

```
naviseccli -h ss1_spa backendbus -resetspeed
```

```
WARNING: the command you are about to execute requires  
that the array 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)?
```

## baseuuid

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

**Description** The **naviseccli baseuuid** 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 storage-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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
baseuuid [uuid-base-value]
```

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

**Options** *uuid-base-value*

Specifies the UUID base.

**Output** If the version of FLARE software running on the SP does not support this command, a `Not Supported` error message is printed to **stderr**.

**Examples** **naviseccli -h ss1\_spa baseuuid**

```
Base UUID of the array: 12345
```

**naviseccli -h ss1\_spa baseuuid 34**

```
Base UUID of the array is set to 34.
```

These commands get, then set the UUID base of the storage system containing the SP **ss1\_spa**.

## bind

### Creates a LUN within an existing RAID group

#### Description

The **naviseccli bind** command binds a LUN within an existing RAID group. You can also create a hot spare RAID group through a single command. You can create a RAID group with the **createrg** command ([page 105](#)).

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

---

**Note:** As soon as a LUN is bound on a CX series storage system running FLARE or Access Logix version 02.04.x.xx.5.xxx or lower, you should start a background verify on the LUN with ASAP priority and set the sniffer rate for the LUN to 5. You can perform both these operations with the **setsniffer** command (see the *Examples* section of the command). You should not send data to the LUN until the background verify operation is complete. The background verify operation will eliminate any latent soft media errors prior to use of the LUN.

As soon as a LUN is bound on a CX series storage system running FLARE or Access Logix software version 02.05.x.xx.5.xxx or later, a background verify runs automatically unless you disabled it with the **-noinitialverify** switch.

---

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



## LUN parameters for different RAID types

RAID type	Bind parameters
RAID 6, RAID 5, RAID 1/0	LUN number Rebuild time Stripe size Read and/or write caching
RAID 3	LUN number Rebuild time
RAID 1	LUN number Rebuild time Read and/or write caching
RAID 0	LUN number Stripe size Read and/or write caching
Disk (individual disk unit)	LUN number Read and/or write caching
Hot Spare	LUN number

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

First - RAID 1/0

Second - RAID 6, RAID 5, RAID 3, and RAID 0

Third - RAID 1

Fourth - Individual units and hot spares

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

### 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, and read caching is recommended for RAID 6.
RAID 5*	You must bind a <i>minimum</i> of 3 disk modules and <i>no more than</i> 16 disk modules. We recommend you bind 5 modules for more efficient use of disk space.

LUN to bind	Restrictions and recommendations
RAID 3	<p>You must bind <i>exactly</i> 5 or 9 disk modules. Select 5 drives for quicker rebuild times or 9 drives for more efficient use of space.</p> <p><b>IMPORTANT:</b></p> <p>For RAID 3, write caching is supported only on CX3 series, CX series storage system running FLARE version 02.06.xxx.x.xxx or higher, or a dual-SP AX4-5 series storage system running Navisphere Manager. For a storage system other than a CX3 series, CX series running FLARE version 02.06.xxx.x.xxx or higher, or a dual-SP AX4-5 series storage system running Navisphere Manager, write-caching is not supported; therefore, when binding RAID 3 LUNs, the <b>-wc</b> <i>cache-flags</i> switch will not apply.</p>
RAID 1	You must bind <i>exactly</i> 2 disk modules.
RAID 0	You must bind a <i>minimum</i> of 3 disk modules, and <i>no more than</i> 16 disk modules.
RAID 1/0	You must bind a <i>minimum</i> of 2 disk modules, and an even number of modules, but <i>no more than</i> 16 modules. Navisphere Manager 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
Hot spare	<p>You <i>cannot</i> bind certain disk modules as a hot spare on a storage system with Fibre Channel disks. For more information about such disk modules, refer to the configuration planning guide for your system.</p> <p>The capacity of a disk module bound as a hot spare must be at least as great as the capacity of the largest disk module that it might replace.</p> <p><b>IMPORTANT</b> The <b>-c</b> <i>cache-flags</i> switch does not apply to a hot spare. When you check status on a hot spare, caching always appears disabled.</p>

- \* You can bind one less module per LUN than you will eventually use. However, the LUN will operate in a degraded mode until a module is installed in the empty slot and the storage system integrates it into the LUN.

You can enable either read or write caching, both read and write caching, or neither type of caching for any type of LUN except for a hot spare. For caching to occur, storage-system caching must also be enabled ([page 87](#) for RAID group storage systems). Whether you should use caching for a specific LUN depends on the RAID type. Recommendations for caching with different RAID types follow.

## Caching recommendations for different RAID types

RAID type	Write caching	Read caching
RAID 0	Acceptable	Recommended
RAID 1	Acceptable	Recommended
RAID 1/0	Acceptable	Recommended
RAID 3*	Recommended	Recommended
RAID 5	Highly recommended	Recommended
RAID 6	Highly recommended	Recommended
Disk	Acceptable	Recommended
Hot spare	Not Allowed	Not Allowed

\* For RAID 3, write caching is supported only on CX3 series or CX series storage system running FLARE version 02.06.xxx.x.xxx or higher, or a dual-SP AX4-5 series storage system running Navisphere Manager. Read caching is only supported on a CX3 series, CX series, or a dual-SP AX4-5 series storage system running Navisphere Manager.

**Note:** LUNs cannot use write caching until you enable the storage-system write cache ([page 283](#)). LUNs cannot use read caching until you enable the read cache for the SP that owns the LUN ([page 87](#) for RAID group storage systems). 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) or stripe size parameters, you must unbind the LUN ([page 324](#)), and then rebind it with the correct parameters. If you need to change the maximum rebuild time, read cache state, or write cache state, use the **chglun** command ([page 93](#)).

**Syntax** `bind` is used with `naviseccli` (described on [page 41](#)) as follows:

```
bind raid-type [lun] -rg rgID
[-aa auto_assignment] [-cap capacity] [-elsz stripe-element-size]
[-noinitialverify][-offset logical-block-address] [-pl placement]
[-r rebuild-priority] [-rc read-cache] [-sp a | b] [-sq size-qualifier]
[-v verify-priority] [-wc write-cache]
```

To create a hot spare RAID group for AX4-5 series and CX series storage systems, use the following syntax:

```
bind raid-type [lun] disks [-aa auto_assignment] [-c cache-type]
[-elsz stripe-element-size] [-n min_latency_reads] [-noinitialverify]
[-sp a | b] [-r rebuild-priority] [-v verify-priority] [-z stripe-count]
```

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

**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

**hs** = hot spare

---

**Note:** RAID 6 is available only for CX series storage systems running FLARE 02.26.xxx.5.yyy or higher, for CX3 series storage systems running FLARE 03.26.xxx.5.yyy or higher, for CX4 series storage systems, and for AX4-5 series storage systems running FLARE 02.23.050.5.7xx or higher.

---

All RAID types are not supported on all storage systems.

*lun*

Specifies a decimal number to assign to the LUN (valid range is shown on [page 46](#)). If 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 ([page 105](#)).

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

**-aa** *auto\_assignment* (CX series only)

Enables or disables auto-assignment 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.)

**-c** *cache-type* (Not supported on AX4-5 series)

You can enable either read or write caching, both read and write caching, or neither type of caching for any LUN (logical unit) except a RAID 3 LUN or hot spare. For caching to occur, storage-system caching must also be enabled. Whether you should use caching for a specific LUN depends on the RAID type.

Recommendations for using caching for the different RAID types, see the Description section.

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

**none** = no caching

**read** = read caching

**write** = write caching

**rw** = read and write caching

The default is read and write caching.

**-elsz** *stripe-element-size* (Not supported on AX4-5 series and AX series)

---

**Note:** Use this option only at the direction of CLARiiON Performance Engineering.

---

**-noinitialverify** (Not supported on AX4-5 series and AX series)

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.

**-offset** *logical-block-address* (Not supported on AX4-5 series and AX series)

Sets the LUN's starting Logical Block Address (LBA) to begin at *an offset location* on the RAID group, such that the user data is aligned with a stripe boundary. Enter the LBA at which user data starts in the host file system and the storage system will compute the correct offset automatically. The default is the lowest available disk address. Use a decimal, not hexadecimal, number. Hexadecimal numbers are not allowed. You can use the **getlun -offset** switch to learn the offset of an existing LUN.

**-pl** *placement* (Not supported on AX4-5 series and AX series)

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

**-r** *rebuild-priority* (Not supported on AX4-5 series and AX series)

Sets the rebuild priority. The priority at which to reconstruct data on either a hot spare or a new disk module that replaces a failed disk module in a LUN. This priority determines the resources that the SP devotes to rebuilding instead of to normal I/O activity. It applies to all RAID LUNs except RAID 0. Valid values are ASAP, High, Medium, and Low. The default is ASAP. A rebuild priority of ASAP rebuilds the disk module as quickly as possible, but degrades response time somewhat. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuild over a longer period of time by specifying the Low rebuild priority.

---

**Note:** You cannot change the rebuild priority for a LUN that is in the process of rebuilding. You can change the value, but it will have no effect.

---

The time it takes to actually rebuild a LUN can vary significantly, especially for a RAID 1/0 LUN, since a RAID 1/0 LUN with  $n$  disk modules can continue functioning with up to as many as  $n/2$  failed drive modules and only one drive at a time is rebuilt.

---

**Note:** The rebuild priority properties are unavailable for RAID 0, disk, or hot spare LUN.

---

**-rc** *read-cache* (AX4-5 series running Navisphere Manager and CX series only)

Enables or disables read cache functionality for this specific LUN.

*read-cache* values:

**0** = Disables read cache

**1** = Enables read cache (the default)

**-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

**sc** = stripe count

**bc** = block count

---

**Note:** AX4-5 series storage systems support only megabyte (MB), gigabyte (GB), and terabyte (TB) size qualifiers.

The stripe count (sc) and block count (bc) size qualifiers are supported only on AX4-5 series storage systems running Navisphere Manager.

---

**-v** *verify-priority* (AX4-5 series running Navisphere Manager and CX series only)

Sets the priority at which the data on each LUN is verified if an SP fails and is replaced during I/O operations. *verify-priority* specifies the priority. Valid values are ASAP, High, Medium, and Low. A lower priority verifies LUN data more slowly and has less impact on overall system performance. The default is High.

---

**Note:** The verify priority properties are unavailable for RAID 0, disk, or hot spare LUN.

---

**-wc** *write-cache* (Dual-SP AX4-5 series running Navisphere Manager only)

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

**0** = Disables write cache

**1** = Enables write cache (the default).

**-z** *stripe-count*

Sets the number of stripes in a LUN. The *stripe-count* specifies the number of stripes. Valid values are any number greater than or equal to 0. The default is 0, which binds the maximum number of stripes available. For example, when binding a RAID 5 LUN with a stripe count of 2, partition the LUN into 2 stripes, thus preventing access to the remainder of the available space. This switch is useful for fast bind operations to produce LUNs for test purposes.

**Output** There is no output. Errors are printed to **stderr**.



**Examples** `naviseccli -h ss1_spa createrg 0 0_0_1 0_0_2 0_0_3  
0_0_4 0_0_5`

This command creates RAID group 0 from the five specified disks (disks 0 through 4 in the DPE enclosure). The new RAID group has no RAID type until it is bound (next).

**naviseccli -h 10.14.46.176 bind hs 66 0\_0\_9**

```
Excerpt of the output of getrg:  
RaidGroup ID: 1  
RaidGroup Type: hot_spare  
RaidGroup State: Valid_luns  
List of disks: Bus 0 Enclosure 0 Disk 9  
List of luns: 66  
Max Number of disks: 1  
Max Number of luns: 1  
Raw Capacity (Blocks): 279701232  
Logical Capacity (Blocks): 279701232  
Free Capacity (Blocks,non-contiguous): 112  
Free contiguous group of unbound segments: 112  
Defrag/Expand priority: Medium  
Percent defragmented: 100  
Percent expanded: 100  
Disk expanding onto: N/A  
Lun Expansion enabled: NO  
Legal RAID types: hot_spare  
Hot Spare Replacing RAID Group: Inactive
```

## cachecard -initialize

### Initializes the cache memory card

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

**Syntax** **cachecard -initialize** is used with **naviseccli** (described on [page 41](#)) as follows:

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

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

**Options** **-o**  
Executes the command without prompting for confirmation.

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

**Example** **naviseccli -h ss1\_spa cachecard -initialize**

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

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

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

## cachecard -list

### Lists the state information of the cache memory card

**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 display.

The possible hardware states are listed in the following table:

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 storage system.

**Syntax** **cachecard -list** is used with **naviseccli** (described on [page 41](#)) as follows:

**cachecard -list [-size] [-hwstate]**

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

**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.

**Output**      Total Memory:      512MB  
                 Hardware State:    Ok

**Example**      `naviseccli -h ss1_spa cachecard -list`

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

## chglun

### Changes parameters on a currently bound LUN

#### 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 on [page 350](#).

#### Caching recommendations for different RAID types

RAID type	Write caching	Read caching
RAID 0	Acceptable	Recommended
RAID 1	Acceptable	Recommended
RAID 1/0	Acceptable	Recommended
RAID 3*	Recommended	Recommended
RAID 5	Highly recommended	Recommended
RAID 6	Highly recommended	Recommended
Disk	Acceptable	Recommended
Hot Spare	Not Allowed	Not Allowed

\* For RAID 3, write caching is supported only on CX3 series, CX series storage system running FLARE version 02.06.xxx.x.xxx or higher, or a dual-SP AX4-5 series storage system running Navisphere Manager. Read caching is supported only on a CX3 series, CX series, or a dual-SP AX4-5 series storage system running Navisphere Manager.

**Note:** LUNs cannot use write caching until you enable the write cache ([page 283](#)). LUNs cannot use read caching until you enable the read cache for the SP that owns the LUN. 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
chglun -l lun
[-a auto-assign] [-c cache-type] [-d default-owner] [-dp] [-i idle-thresh]
[-mp max-prefetch-blocks] [-name newname]
[-pc max-prefetch-IO] [-pd min-prefetch-size] [-pf prefetch-type]
[-pm pf-size/mult] [-r rebuild-priority] [-rt retain-prefetch-data]
[-sm prefetch-segmnt-size/mult] [-t idle-delay-time] [-v verify-priority]
[-w write-aside]
```

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

**Options** -l *lun*

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

-a *auto-assign* (CX series only)

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 auto-assignment for specified LUN

1 = Enables auto-assignment for specified LUN

-c *cache-type* (AX4-5 series running Navisphere Manager and CX series only)

Enables either read or write caching, both read and write caching, or neither type of caching for the LUN. For caching to occur, storage-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 83](#).

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

**none** = no caching

**read** = read caching

**write** = write 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

**-dp** (AX4-5 series running Navisphere Manager and CX series only)

Sets the default prefetch values.

**-i** *idle-thresh* (Not supported on AX4-5 series and AX series)

Sets the maximum number of I/Os that can be outstanding to a LUN and still have the LUN be considered idle. *idle-thresh* specifies the maximum number of I/Os. Used to determine cache flush start time. Legal *idle-thresh* values range from 0 through 254.

**-mp** *max-prefetch-blocks* (AX4-5 series running Navisphere Manager and CX series only)

Sets the maximum number of disk blocks to prefetch in response to a read request. *max-blocks* specifies the maximum number.

**-name** *newname*

Sets the LUN name. *newname* specifies the new name for the LUN. This switch applies only to storage systems running a version of FLARE 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 storage systems.

**-pc** *max-prefetch-IO* (AX4-5 series running Navisphere Manager and CX series only)

Sets the maximum number of outstanding I/Os that the LUN can have and still perform prefetching. *max-IO* specifies the maximum number of I/Os. Valid values are 0 through 100.

**-pd** *min-prefetch-size* (AX4-5 series running Navisphere Manager and CX series only)

Sets the prefetch disable size. *min-size* specifies, in blocks, the size of the smallest read request for which prefetching is disabled. Prefetching is disabled for all read requests whose size is equal to or larger than *min-size*. Valid values are 0 through 65534.

**-pf** *prefetch-type* (Not supported on AX4-5 series and AX series)

Sets or disables the prefetch type as follows:

0 = Disables prefetch

1 = Constant length prefetch type

2 = Variable length prefetch type

**-pm** *pf-size-or-multiplier* (AX4-5 series running Navisphere Manager and CX series only)

Sets the amount of data or the prefetch multiplier prefetched for one host read request as follows:

For constant length prefetching: 0-2048.

For variable length prefetching: 0-32

If *prefetch-type* (see below) is set to constant length, *pf-size-or-multiplier* specifies the number of blocks to prefetch. The valid range is 0-2048 blocks.

If *prefetch-type* is set to variable length, *pf-size-or-multiplier* specifies the multiplier used to determine the amount of data to prefetch. The valid range is 0-32. For example, if the value is 8, the amount of data to prefetch is 8 times the amount of data requested.

The default value for *pf-size-or-multiplier* is 4.

**-r** *rebuild-priority* (Not supported on AX4-5 series and AX series)

Sets the rebuild priority. *rebuild-priority* specifies the priority at which to reconstruct data on either a hot spare or a new disk module that replaces a failed disk module in a LUN. It applies to all RAID LUNs types except RAID 0 and individual disks. Valid values are ASAP, High, Medium, and Low. The default is ASAP. The priority that you specify determines the amount of resource that the SP devotes to rebuilding instead of to normal I/O activity. The default value is adequate for most situations. A rebuild priority of ASAP rebuilds the disk module as quickly as possible, but degrades response time significantly. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuild over a longer period of time by specifying the Low rebuild priority.



---

**Note:** It is not possible to change the rebuild priority for a LUN that is in the process of rebuilding. You can change the value, but it will have no effect.

---

The time it takes to actually rebuild a LUN can vary significantly, especially for a RAID 1/0 LUN, since a RAID 1/0 LUN with  $n$  disk modules can continue functioning with up to as many as  $n/2$  failed drive modules and only one drive at a time is rebuilt.

**-rt** *retain-prefetch-data* (Not supported on AX4-5 series and AX series)

Enables or disables retaining prefetch data functionality.

*retain-prefetch-data* specifies functionality as follows:

**0** = Disables retaining prefetch data functionality

**1** = Enables retaining prefetch data functionality

**-sm** *prefetch-segmnt-size/mult* (AX4-5 series running Navisphere Manager and CX series only)

Sets the segment size for data prefetched for one host read request as follows:

If **-pf** *prefetch-type* is set to constant length, *segmnt-size/mult* specifies the size of the segment in blocks. Valid range: 0-2048 blocks.

If **-pf** *prefetch-type* is set to variable length, *segmnt-size/mult* specifies the multiplier used to determine the segment size. Valid range: 0-32.

**-t** *idle-delay-time* (AX4-5 series running Navisphere Manager and CX series only)

Sets the amount of time that a LUN must be below the *idle-thresh* (see below) in order to be considered idle. *idle-delay-time* specifies the delay time in 100 ms units. For example, 5 equals 500 ms.

Once a LUN is considered idle, any dirty pages in the cache can begin flushing. Legal *idle-delay-time* values range from 0 through 254.

**-v** *verify-priority* (AX4-5 series running Navisphere Manager and CX series only)

Sets the priority at which to verify data on a LUN if an SP fails and is replaced during I/O operations. *verify-priority* specifies priority level. Valid values are ASAP, High, Medium, and Low.

The lower the priority, the slower the LUN is rebuilt and the lesser the impact on overall system performance. The default is Low.

**-w** *write-aside* (Dual-SP AX4-5 series running Navisphere Manager 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 16 through 65534.

---

**Note:** Attempting to use this switch while not in Engineering Mode will result in an error message.

---

**Output** There is no output. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa chglun -l 3 -c write -r high`  
Change LUN 3 to enable write caching and rebuild with high priority.

## chgrg

### Changes properties of a RAID group

**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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
chgrg rgID [-defrag] [-expand disks] [-lex expansion]
[-powersavings on | off] [-pri priority] [-rm yes | no] [-trespass sp]
```

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

**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 storage system.

#### **-defrag**

Starts defragmenting the specified RAID group.

#### **-expand *disks***

Expands the RAID group *rgID* onto the disk(s) specified at a given priority, and may expand LUNs.

*disks* specifies the physical disks to bind. The order in which you specify the disks has no significance. See “[LUN IDs, unique IDs, and disk IDs](#)” on [page 46](#) to verify the format for specifying disks.

**-lex** *expansion* (Not supported on AX4-5 series and AX series)

Available only when the **-expand disks** switch is used. Not available for storage systems that support metaLUNs. *expansion* specifies functionality as follows:

**yes** = Enables LUN expansion for the specified RAID group  
**no** = Disables LUN expansion for the specified RAID group (the default)

**-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.

**-pri** *priority* (Not supported on AX4-5 series and AX series)

Sets the priority for defragmenting or expanding the specified RAID group. *priority* specifies the process priority as follows:

**high** = high priority  
**medium** = medium priority  
**low** = low priority

**-rm yes | no** (Not supported on AX4-5 series and AX series)

Enables or disables the removal of the specified RAID group after the last LUN in it has been unbound (**unbind** command).

*remove* specifies the functionality as follows:

**yes** = Removes the RAID group  
**no** = Does not remove the RAID group (the default)

**-trespass** *sp* (Not supported on CX series, AX4-5 series, and AX series)

Trespasses all LUNs in the RAID group to the specified storage processor. *sp* specifies the SP as follows:

**a** = SP A  
**b** = SP B

**Example** `naviseccli -h ss1_spa chrg 2 -trespass a`

This command trespasses all LUNs in RAID group 2 to SP A.

---

## clearlog

### Clears the SP's error log

- 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` is used with `naviseccli` (described on [page 41](#)) as follows:  
`clearlog`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Output** Confirmation message.
- Example** `naviseccli -h ss1_spa clearlog`  
`unsolicited error log cleared`  
This command deletes the SP logs of the storage system.

## clearstats

### Resets the SP's statistics logging

**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** is used with **naviseccli** (described on [page 41](#)) as follows:  
**clearstats**

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

**Output** Confirmation message.

**Example** **naviseccli -h ss1\_spa clearstats**  
Statistics successfully cleared, statistics logging disabled.

## **-cmdtime**

**Displays the date and endtime of command execution.**

**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.

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

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

**Options** **-cmdtime**  
Displays the date and endtime of command execution.

**Example** **naviseccli -h ss1\_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
```

## copytohot spare

**Copies data from a failing disk onto an existing hot spare**

**Description** This **naviseccli** command initiates the copying of data from a failing disk to an existing hot spare while the original disk remains functioning.

Once the copy is made, the candidate will be faulted and the hot spare activated. When the faulted disk is replaced, the replacement will be copied back from the hot spare.

**Syntax** **copytohot spare** is used with **naviseccli** (described on [page 41](#)) as follows:

**copytohot spare** *disk-position* **-initiate [-0]**

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

**Options** *disk-position*

This specifies the physical disk. See “[LUN IDs, unique IDs, and disk IDs](#)” on [page 46](#) to verify the format for specifying disks.

**-initiate**

Requires the disk number of the proactive candidate, which you must specify.

Optional switch:

**-0**

Executes the command without prompting for confirmation.

**Output** The disk 0\_0\_13 is selected for hot sparing. Its serial number is KFDJ89. The proactive sparing process cannot be aborted and may take a long time to complete. This disk should not be replaced till its status is changed to Faulted at the end of the process. Do you wish to continue?

**Example** **naviseccli -h ss1\_spa copytohot spare 0\_0\_13 -initiate**



## createrg

### Creates a RAID group

**Description** The `naviseccli createrg` command creates a RAID group on the specified disks. When you create a RAID group on a storage system running Navisphere Manager, it has no RAID type. The new group is assigned the RAID type of the first LUN you bind on the group.

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

---

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

---

**Syntax** `createrg` is used with `naviseccli` (described on [page 41](#)) as follows:  
`createrg rgID disks [-pri priority] [-powersavings on | off] [-rm yes | no] [-raidtype r6 | r5 | r3 | r1_0]`

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

**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 storage 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 such an order as to determine 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.

See “LUN IDs, unique IDs, and disk IDs” on [page 46](#) to verify the format for specifying disks.

---

**Note:** The disks in a RAID group must be either all Fibre Channel disks or all ATA disks.

---

**-raidtype r6 | r5 | r3 | r1\_0**

Sets the RAID type of the RAID group during the time of creation. The supported RAID types are r6, r5, r3, and r1\_0. For AX4-5 series storage systems, the RAID6 is supported only on FLARE version 02.23.050.5.7xx or later.

**IMPORTANT**

**The -raidtype is a mandatory switch for AX4-5 series storage systems running Navisphere Express. The -raidtype is an optional switch for AX4-5 series running Navisphere Manager and CX series running FLARE version 03.26.XXX or later. This switch is not supported on AX series storage systems or CX series running FLARE version prior to 03.26.XXX.**

You can create a hot spare RAID group for AX4-5 series and CX series systems using the disk **bind** command described on [page 80](#). You cannot create a hot spare directly for AX4-5 series systems running Navisphere Express. To create a hot spare for AX4-5 series systems running Navisphere Express use disk bind command described on [page 80](#).

**-pri *priority***

Sets the priority for expanding or defragmenting the specified RAID group. *priority* specifies the process priority as follows:

**high**  
**medium**  
**low** (the default)

**-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.

**-rm *yes* | *no***

Enables or disables the removal of the specified RAID group after the last LUN in it has been unbound (**unbind** command).

*remove* specifies the functionality as follows:

**yes** = Removes the RAID group  
**no** = Does not remove the RAID group (the default)

**Example** Navisphere Manager:

```
naviseccli -h ss1_spa createrg 4 0_1 0_2 0_3 0_4 0_5 -pri high
```

This command creates RAID group 4 from the five specified disks, and sets the priority for the expansion/defragmenting processes to high.

Navisphere Express:

```
naviseccli -h ss1_spa createrg 1 0_5 0_6 0_7 0_8 0_9 -pri high  
-rm yes -raidtype r3
```

This command creates RAID group 1 from the five specified disks and removes the specified RAID group after the last LUN in it has been unbound. This command also sets raidtype r3 to the RAID group and the priority for the expansion/defragmenting processes to high.

## emconfiguration

### Configures event monitor

**Description** The **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** is used with **naviseccli** (described on [page 41](#)) as follows:

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

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

**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 ss1\_spa emconfiguration -reloadconfig -disableresponse *time* -enableresponse**

## failovermode

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

**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 if suggested in the *Installation Roadmap for CLARiiON Storage Systems*, which is on the EMC Powerlink website.

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 storage systems without storage group only. For storage systems with storage group, use the command **storagegroup sethost -failovermode** (see [page 305](#)).

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 storage system inaccessible.**

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

**Syntax** **failovermode** is used with **naviseccli** (described on [page 41](#)) as follows:

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

For information about when to use these failover modes, refer to the *Installation Roadmap for CLARiiON Storage Systems*, which is on the EMC Powerlink website.

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

**Options** **-o**

Executes the command without prompting for confirmation.

**Output** If the version of FLARE 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
```

**Example** `naviseccli -h ss1_spa failovermode 0`

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

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

## faults -list

**Lists status information for faulted components on the storage system**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:  
`faults -list`

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

**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:
Asynchronous mirror is faulted
```

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

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

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





## firmware

### Updates the disk firmware

**Description** The `naviseccli firmware` command updates the disk firmware. To update an SP's FLARE software, use the `ndu` command. To update disk firmware, use the `firmware` command as described here.

**Note:** After downloading firmware, the CLI prompts you if a reboot is necessary. In general, downloaded firmware is updated only after the storage system reboots. However, depending on what is included in the firmware file a reboot may or may not be necessary.

**Syntax** `firmware` is used with `naviseccli` (described on [page 41](#)) as follows:  
`firmware filename [-c] [-d disks] [-w] [-o]`

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

**Options** *filename*

Specifies the full pathname of the new firmware image on the destination server. This image contains the microcode. If the full pathname is not entered, the current working directory is added to the beginning of the filename.

**-c**

Specifies a client side download. Indicates that the firmware file resides on a client (the default is no). For example, if you are running the CLI from a machine that is not the storage system server, you can still download firmware if you specify this flag.

**-d disks**

Downloads disk vendor firmware to specified disks.

*disks* specifies the physical disks whose firmware will be updated. The order in which you specify the disks has no significance. See [“LUN IDs, unique IDs, and disk IDs” on page 46](#) to verify the format for specifying disks.

**-o**

Executes the command without prompting for confirmation.

**-w**

Initiates a warm reboot after the firmware is downloaded. A warm reboot performs the same functions as a reboot except that the powerup diagnostics are not executed. In systems that use caching, the caches are re-initialized. A warm reboot executes in under 50 seconds.

**Output** A warning message appears before the command is issued to the storage system, if a reboot is necessary. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa firmware /tmp/mcode.rev -o`  
Downloads **mcode.rev**, which is located in the **tmp** directory.

## flash -ioport

Flashes the LEDs on I/O ports.

**Description** The **naviseccli flash -ioport** command flashes the I/O ports on the CX4 series storage system. Only the I/O ports that are owned by the targeted SP are flashed.

**Syntax** **flash -ioport** is used with **naviseccli** (described on [page 41](#)) as follows:

```
flash -ioport [-fe logicalportID | -be logicalportID |
-pslotid physicalslotID -portid physicalportID | -byusage mirrorview]
[on | off]
```

**Prerequisites** For **naviseccli**, you must have a user account on the storage system on where want to execute the command.

**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** *physicalslotID* **-portid** *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.

**-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.

**Examples**

To turn flashing **on** for a specified physical port:

```
naviseccli -h 10.14.125.176 flash -ioport -pslotid 0 -pportid 3 on  
Slot 0, Physical Port 3 (Uninitialized): ON
```

To turn flashing **on** by specifying a logical port ID:

```
naviseccli -h 10.14.125.176 flash -ioport -be 3 on  
Slot 3, Physical Port 1 (BE 3): ON
```

To turn flashing **on** by specifying the usage of I/O port:

```
naviseccli -h 10.14.125.176 flash -ioport -byusage mirrorview on  
Slot 1, Physical Port 3 (FE 1): ON
```

## flash -iomodule

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

**Description** The `naviseccli flash -iomodule` command turns the flashing LEDs on/off for all the I/O ports on a specified module on the CX4 series storage system.

**Syntax** `flash -iomodule` is used with `naviseccli` (described on [page 41](#)) as follows:

`flash -iomodule [slotid] [on | off]`

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

**Options** `slotid`

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

`on | off`

Indicates whether flashing is turned on or off.

**Examples** `naviseccli -h 10.14.5.230 flash -iomodule 0 on`

```
I/O Module 0:ON
```

`naviseccli -h 10.14.5.230 flash -iomodule 0 off`

```
I/O Module 0:OFF
```

## flashleds

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

**Description** The `naviseccli flashleds` 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` is used with `naviseccli` (described on [page 41](#)) as follows:  
`flashleds -e enclosure-number [-b bus-number] [on | off]`

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

**Options**

- `-e enclosure-number`  
 Specifies the disk-array enclosure. In each storage 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 storage systems) or ascending from 0 (for other storage-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.

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

**Examples** `naviseccli -h ss1_sp1 flashleds -e 1 -b 0 on`

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

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

## getagent

### Gets device names and configurations

**Description** The **naviseccli getagent** command returns the names of devices that the destination Navisphere agent controls. The command also displays descriptions of each device with which the Navisphere 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 Navisphere Host Agent is supported only on ESX Server.

**Syntax** **getagent** is used with **naviseccli** (described on [page 41](#)) as follows:  
**getagent** [-cabinet] [-desc] [-mem] [-model] [ -name] [ -node] [-os]  
 [-peersig] [-prom ] [-rev] [-scsiid] [-serial] [-sig] [-spid] [-type] [-ver]

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

**Options** **-cabinet**

Returns a description of the cabinet type.

**-desc**

Returns an ASCII description of the device (taken from the Navisphere agent configuration file).

**-mem**

Returns, in MB, the size of the SP's memory. Each SP in the storage 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 Navisphere 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 storage system should be running the same revision of PROM code. PROM code is updated automatically when you update an SP's FLARE software and choose the reboot option.

**-rev**

Returns the microcode revision number of the FLARE software that the SP is running. Each SP in the storage system will be running the same revision of FLARE software. You can update an SP's FLARE 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 storage-system installation and maintenance manual.

**-serial**

Returns the serial number of the storage 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.

**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 storage system, the output may vary depending upon whether or not the host is in a storage group.

**Output from an SP agent**

```
Agent Rev:      6.1.0 (8.4)
Name:          K10
Desc:
Node:          A-WRE00021400494
Physical Node: K10
Signature:     588114
Peer Signature: 588069
Revision:      2.01.1.01.4.001
SCSI Id:       0
Model:         600
Model Type:    Rackmount
Prom Rev:      2.12.00
SP Memory:     2048
Serial No:     WRE00021400494
SP Identifier: A
Cabinet:      xPE
```

**Example** This example issues the **getagent** command to storage system `ss1_spa`.

```
naviseccli -h ss1_spa getagent
```

Sample output follows.

## getall

### Gets comprehensive storage-system information

#### Description

The **naviseccli getall** command returns an extensive list of storage-system information, including:

- ◆ Agent and host information
- ◆ Storage-system (array) 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 Navisphere release. The amount of this information may change (usually it increases) with each new Navisphere 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** is used with **naviseccli** (described on [page 41](#)) as follows:

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

#### Prerequisites

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

**Options -host**

Displays only information about the host and the agent that manage the storage system. This is a composite of information displayed by the **getagent**, **remoteconfig**, and **ndu** commands.

**-array**

Displays only information about the storage 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 CX4 series storage 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 a composite of information the **getcache** and **getconfig** commands.

**-disk**

Displays information about disks and enclosures in the storage system and new disk spin-down statistics. This is the same information displayed the **getdisk** command displays.

**-drivetype**

Displays the drive type of the LUN.

**-lun**

Displays only information about each logical unit (host and storage system) in the storage 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.

**Output formats     Agent/Host information (same as returned by command getsp)**


---

Server IP Address:	10.15.22.176	Server/SP IP address
Agent Rev:	5.3.0 (5.0)	Revision of the agent.
Agent/Host Information		
-----		
Desc:	MY_K10	ASCII description of device.
Node:	A-F20011000207	Device name or location in file system of device. Use this value when specifying the command line device name.
Physical Node:	K10	
Signature:	196618	Unique signature of SP
Peer Signature:	176166	Unique signature of peer SP
SCSI Id:	0	SCSI id of SP host connection
SP Identifier:	A	SP Identifier (A or B)
Model:	4700-2	Model number of SP
Model Type:	Rackmount	Model Type (deskside, rackmount)
Prom Rev:	5.94.00	Prom revision.

---

---

SP Memory:	817	SP memory in megabytes.
Serial No:	F20011000207	Serial number of device.
Cabinet:	DPE	cabinet type (DPE)

**Installed packages:**

Name of the software package:Base	Package name
Revision of the software package: 01_89_35	Package revision
Commit Required:	NO YES / NO
Revert Possible:	NO YES / NO
Active State:	YES YES / NO
Dependent packages: MirrorView 01_89_25, Navisphere 5.3.0	Dependency information
Required packages: Navisphere >5.2.0.0.0, Base <01_89_99, Base >01_80_00,	Required package revision information
Is installation completed: YES	YES / NO
Is this System Software: NO	YES / NO

---

**Storage-system information (-array switch) (same information is  
returned by commands getcrus and initializearray -list).**

---

**Array Information**

-----

Array Name: Array_2	
Array UID:	
50:06:01:60:20:04:A0:CF	
Array Name: Storage_2	Storage-system name
Array UID:50:06:01:60:20:04:A0:CF	Storage-system 16-byte unique identification number

---

```

DAE/SP Bus 0 Enclosure 0
Bus 0 Enclosure 0 Fan A State: Present
Bus 0 Enclosure 0 Fan B State: Present
Bus 0 Enclosure 0 Power A State: Present
Bus 0 Enclosure 0 Power B State: Present
SP A State: Present
SP B State: Present
Bus 0 Enclosure 0 SPS A State: Present
Bus 0 Enclosure 0 SPS B State: Present
Bus 0 Enclosure 0 LCC A State: Present
Bus 0 Enclosure 0 LCC B State: Present
Bus 0 Enclosure 0 LCC A Revision: 3
Bus 0 Enclosure 0 LCC B Revision: 3
Bus 0 Enclosure 0 LCC A Serial #: F4001020156
Bus 0 Enclosure 0 LCC B Serial #: N/A
DAE Bus 0 Enclosure 1
Bus 0 Enclosure 1 Fan A State: Present
Bus 0 Enclosure 1 Power A State: Present
Bus 0 Enclosure 1 Power B State: Present
Bus 0 Enclosure 1 LCC A State: Present
Bus 0 Enclosure 1 LCC B State: Present
Bus 0 Enclosure 1 LCC A Revision: 3
Bus 0 Enclosure 1 LCC B Revision: 0
Bus 0 Enclosure 1 LCC A Serial #: F4001080137
Bus 0 Enclosure 1 LCC B Serial #: N/A

```

**(The CRU display continues for each enclosure on bus 0, then for each enclosure on bus 1)**

```

Raid Group ID: 223
Logical Units Used: 223
List of disks: Bus 0
Enclosure 0 Disk 2
    Bus 0 Enclosure 0 Disk 3
    Bus 0 Enclosure 0 Disk 4
    Bus 0 Enclosure 0 Disk 5
    Bus 0 Enclosure 0 Disk 6

```

**HBA or Switch information**

```

HBA Information
-----
Information about each HBA:
HBA UID:
01:02:03:04:05:06:07:08:09:0A:0B:0C:0D:0E:01:12
Server Name: navi2280.us.dg.com
Server IP Address: 12.34.56.78
HBA Model Description:
HBA Vendor Description:
HBA Device Driver Name:

```

Customer replaceable unit (CRU) information

Information about the PSM LUN (created at system installation)

HBA 32-byte unique identification number

Server name.

Server IP address.

If any

If any

If any

---

 Information about each port of this HBA:

SP Name:	SP A	SP (A or B)
SP Port ID:	0	Can be 0 through 3.
HBA Devicename:	N/A	Name in file system if available
Trusted:	NO	YES or NO
Logged In:	NO	YES or NO
Defined:	YES	YES or NO
Initiator Type:	0	Can be 0 through 3
StorageGroup Name:	Server1_SG	Storage group(s) on this port

**(HBA and HBA port listings continue here.)**

## Information about each SPPORT:

SP Name:	SP A	A or B
SP Port ID:	0	0 to 3
SP UID:	50:06:01:60:20:04:A0:CF:50:06:01:60:40:04:A0:CF	32-byte unique ID number of SP
Link Status:	Down	Up or Down
Port Status:	DISABLED	Enabled or Disabled
Switch Present:	NO	YES or NO
ALPA Value:	0	0 to 255
Speed Value:	1	Speed of 1 or 2 Gbs

**(SP A and SP B port information appears here.)**

---

 SP Information
 

---

```

-----
Storage Processor:                SP A                A or B
Storage Processor Network Name:   naviqa22176      Network information
Storage Processor IP Address:     10.15.22.176
Storage Processor Subnet Mask:    255.255.255.0
Storage Processor Gateway Address: 10.15.22.1
System Fault LED:                 OFF              Status of fault LED.
Statistics Logging:               ON               Serve for performance
                                                         monitoring
System Date:                      08/15/2001      SP clock information
Day of the week:                  Wednesday
System Time:                      08:37:
Max Requests:                     1               Performance
Average Requests:                 4294967295      information (recorded
Prct Busy:                         0               only if Statistics
Prct Idle:                         0               logging is on)
Hard errors:                       0
Read_requests:                     0
Write_requests:                    0
Total Reads:                       0
Total Writes:                      0
Blocks_read:                       0
Blocks_written:                    0
Sum_queue_lengths_by_arrivals:     0
Arrivals_to_non_zero_queue:        0
Hw_flush_on:                       0
Idle_flush_on:                     31
Lw_flush_off:                       0
Write_cache_flushes:               0
Write_cache_blocks_flushed:        0
Internal bus 1 busy ticks:          4294967295
Internal bus 1 idle ticks:          4294967295
Internal bus 2 busy ticks:          0
Internal bus 2 idle ticks:          0
Internal bus 3 busy ticks:          0
Internal bus 3 idle ticks:          0
Internal bus 4 busy ticks:          0
Internal bus 4 idle ticks:          0
Internal bus 5 busy ticks:          0
Internal bus 5 idle ticks:          0
Controller busy ticks:              20787
Controller idle ticks:              1617657
  
```

**(SP B information appears here.)**

---



**The physical slot and physical port information for the SP ports that the getall command with the -hba switch returns.**

Information about each SPPORT:

```

SP Name:                SP A
SP Port ID:             0
SP UID:
50:06:01:60:B9:A0:13:86:50:06:01:60:39:A0:13:86
Link Status:           Up
Port Status:           Online
Switch Present:        YES
Switch UID:
10:00:08:00:88:E3:0A:EF:20:1D:08:00:88:E3:0A:EF
SP Source ID:          7871763
ALPA Value:            0
Speed Value :          4Gbps
Auto Negotiable :      NO
Available Speeds:
1Gbps
2Gbps
4Gbps
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 :     1
Physical Port ID :    3
Usage:                 General

```

```

SP Name:                SP A
SP Port ID:             1
SP UID:
50:06:01:60:B9:A0:13:86:50:06:01:61:39:A0:13:86
Link Status:           Down
Port Status:           DISABLED
Switch Present:        NO
ALPA Value:            0
Speed Value :          N/A
Auto Negotiable :      YES
Available Speeds:     1Gbps
2Gbps
4Gbps
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 :2
Usage:              General
```

**Management port information returned by the command `getall -sp` (The information is the same as returned by command `networkadmin -get -all`).**

#### SP Information

```
-----
Storage Processor: SP A
Storage Processor Network Name: OEM-2MYA58ZGAQT
Storage Processor IP Address: 10.4.78.158
Storage Processor Subnet Mask: 255.255.255.0
Storage Processor Gateway Address: 10.4.78.1
```

#### Management Port Settings:

```
Current Speed: 100 Mbps/half duplex
Requested Speed: 100 Mbps/half duplex
Link Status: Link-up
Auto-Negotiate: No
Capable Speeds:
10 Mbps half/full duplex
100 Mbps half/full duplex
1000 Mbps half/full duplex
Auto
```

**Sample output displaying resume information that the getall  
-resume command for the SPE enclosure on a CX4 series storage  
system returns.**

---

```

Server IP Address:      10.14.20.175
Agent Rev:             6.28.0 (1.38)

Resume Information
-----
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:          N/A
  Programmable Revision:     N/A
  EMC Product Serial Number:
XXXXXXXXXXXXXXXXX
  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: N/A
Programmable Revision: N/A
```

```
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: LONGHUATOWN, SHENZHEN,
CHINA
Year of Manufacture: 2006
Month of Manufacture: 11
Day of Manufacture: 09
Assembly Name: 4GB FCDAE STILETTO W/ FLANGES
Programmable Name: N/A
Programmable Revision: N/A
```

---

**SP cache information (same as returned by commands getcache and getcontrol).**

---

Cache Information

-----

System Buffer (spA):	58 MB
System Buffer (spB):	58 MB
SP Read Cache State	Enabled
SP Write Cache State	Enabled
Cache Page size:	2
Write Cache Mirrored:	YES
Low Watermark:	40
High Watermark:	60
SPA Cache pages:	60657
SPB Cache pages:	0
Unassigned Cache Pages:	0
Read Hit Ratio:	29
Write Hit Ratio:	0
Prct Dirty Cache Pages =	0
Prct Cache Pages Owned =	49
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:	21:00
SPA Physical Memory Size =	817
SPA Free Memory Size =	399
SPB Physical Memory Size =	817
SPB Free Memory Size =	409
SPA Read Cache Size =	110
SPA Write Cache Size =	250
SPA Optimized Raid 3 Memory Size =	0
SPB Read Cache Size =	100
SPB Write Cache Size =	250
SPB Optimized Raid 3 Memory Size =	0
HA Cache Vault=	Not Available

---

**Disk information (same as getdisk returns) and HBA information.****naviseccli -h 10.14.46.236 getall -disk**

```
Server IP Address:      10.14.46.236
Agent Rev:             6.30.0 (1.34)
```

## All Disks Information

-----

```
Bus 2 Enclosure 0 Disk 0
Vendor Id:             SEAGATE
Product Id:           STT34007 CLAR400
Product Revision:     C003
Lun:                  Unbound
Type:                 N/A
State:                Unbound
Hot Spare:            NO
Prct Rebuilt:         Unbound
Prct Bound:           Unbound
Serial Number:        3RJ0QQPS
Sectors:              0 (0)
Capacity:             375600
Private:              Unbound
Bind Signature:       0x0, 0, 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:       This disk does not belong to a RAIDGroup
Clariion Part Number: DG118032576
Request Service Time: N/A
Read Requests:       0
Write Requests:      0
Kbytes Read:         0
Kbytes Written:      0
Stripe Boundary Crossing: None
Drive Type:          Fibre Channel
Clariion TLA Part Number: 005048775
User Capacity:       0
Idle Ticks:          157286
Busy Ticks:          0
Current Speed:       4Gbps
Maximum Speed:       4Gbps
Queue Max:           N/A
Queue Avg:           N/A
```

```

Prct Idle:          1
Prct Busy:          0
Hardware Power Savings Qualified: NO
Hardware Power Savings Eligible: NO
Power Savings State: Full Power
Current Power Savings Log Timestamp: 06/24/09 12:48:23
Spinning Ticks: 0
Standby Ticks: 0
Number of Spin Ups: 0

```

**(Information appears here for additional disks on bus in this enclosure, then on other bus 0 enclosures, then on bus 1 and each bus 1 enclosure.)**

---

### LUN information (same information as returned by command getlun)

All logical Units Information

```

-----
LOGICAL UNIT NUMBER 0
Name                      LUN 0
UID:
60:06:01:F9:02:6C:00:00:CC:F3:8F:5F:A5:8D:
D5:11
Is Private:                NO
Snapshots List:            None
MirrorView Name if any:    Not Mirrored
Dual Simultaneous Access:  N/A
Prefetch size (blocks) =   0
Prefetch multiplier =      4
Segment size (blocks) =    0
Segment multiplier =       4
Maximum prefetch (blocks) = 512
Prefetch Disable Size (blocks) = 129
Prefetch idle count =     40
Variable length prefetching YES
Prefetched data retained   YES

```

---

Read cache configured according to specified parameters.

```

Bus 0 Enclosure 0 Disk 0 Enabled
Bus 0 Enclosure 0 Disk 1 Enabled
Bus 0 Enclosure 0 Disk 7 Enabled
Bus 0 Enclosure 0 Disk 8 Enabled
Bus 0 Enclosure 0 Disk 9 Enabled

```

---

Read cache information

---

Total Hard Errors:	0	Hard read and write errors
Total Soft Errors:	0	Soft (correctable) errors
Total Queue Length:	0	
RAID-3 Memory Allocation:	0	
Minimum latency reads	N/A	
Read Histogram[0]	0	Performance information (recorded only if statistics logging is on)
Read Histogram[1]	0	
<b>(Read Histogram 2 through 10 information appears.)</b>		
Write Histogram[0]	0	
Write Histogram[1]	0	
<b>(Write Histogram 2 through 10 information appears.)</b>		
Read Requests:	0	
Write Requests:	0	
Blocks read:	0	
Blocks written:	0	
Read cache hits:	0	
Read cache misses:	0	
Prefetched blocks:	0	
Unused prefetched blocks:	0	
Write cache hits:	0	
Forced flushes:	0	
Read Hit Ratio:	0	
Write Hit Ratio:	0	
RAID Type:	RAID5	RAID type
RAID Group ID:	0	Group ID
State:	Bound	State

---



---

```
Stripe Crossing:          0
Element Size:            128
Current owner:           SP B
Offset:                  0
Auto-trespass:           DISABLED
Auto-assign:             DISABLED
Write cache:             ENABLED
Read cache:              ENABLED
Idle Threshold:          0
Idle Delay Time:         20
Write Aside Size:        1023
Default Owner:           SP B
Rebuild Priority:         ASAP
Verify Priority:          Low
Prct Reads Forced Flushed: 0
Prct Writes Forced Flushed: 0
Prct Rebuilt:            100
Prct Bound:              100
LUN Capacity(Megabytes): 1024
LUN Capacity(Blocks):    2097152
```

**(Listing for additional LUNs appears here.)**

---

---

```
LOGICAL UNIT NUMBER 1018
Prefetch size (blocks) =          0
Prefetch multiplier =             4
Segment size (blocks) =          0
Segment multiplier =             4
Maximum prefetch (blocks) =     4096
Prefetch Disable Size (blocks) = 4097
Prefetch idle count =            40
```

```
Variable length prefetching YES
Prefetched data retained    YES
```

```
Read cache configured according to
specified parameters.
```

```
Total Hard Errors:          0
Total Soft Errors:          0
Total Queue Length:        68
Name                        LUN 49
Minimum latency reads N/A
```

```
Read Histogram[0] 0
Read Histogram[1] 0
Read Histogram[2] 0
Read Histogram[3] 0
Read Histogram[4] 0
Read Histogram[5] 0
Read Histogram[6] 0
Read Histogram[7] 0
Read Histogram[8] 0
Read Histogram[9] 0
Read Histogram overflows 0
```

```
Write Histogram[0] 0
Write Histogram[1] 0
Write Histogram[2] 0
Write Histogram[3] 0
Write Histogram[4] 0
Write Histogram[5] 0
Write Histogram[6] 0
Write Histogram[7] 0
Write Histogram[8] 0
Write Histogram[9] 0
Write Histogram overflows 0
```

---

---

```

Read Requests:          0
Write Requests:        0
Blocks read:           0
Blocks written:        0
Read cache hits:       0
Read cache misses:     N/A
Prefetched blocks:     0
Unused prefetched blocks: 0
Write cache hits:      0
Forced flushes:        0
Read Hit Ratio:        N/A
Write Hit Ratio:       N/A
RAID Type:             RAID5
RAIDGroup ID:          5
State:                 Faulted
Stripe Crossing:      0
Element Size:          128
Current owner:         Unknown
Offset:                0
Auto-trespass:         ENABLED
Auto-assign:           ENABLED
Write cache:           ENABLED
Read cache:            ENABLED
Idle Threshold:        0
Idle Delay Time:       20
Write Aside Size:      2048
Default Owner:         SP A
Rebuild Priority:      ASAP
Verify Priority:        ASAP
Prct Reads Forced Flushed: 0
Prct Writes Forced Flushed: 0
Prct Rebuilt:          100
Prct Bound:            100
LUN Capacity(Megabytes): 1024
LUN Capacity(Blocks):  2097152
UID:
60:06:01:60:00:D0:16:00:F2:A2:8C:8D:FB:C
7:D6:11
Blocks Read SPA:       0
Blocks Read SPB:       0
Blocks Written SPA:    0
Blocks Written SPB:    0
Read Requests SPA:     0
Read Requests SPB:     0
Write Requests SPA:    0
Write Requests SPB:    0
LUN Busy Ticks SPA:    0
LUN Busy Ticks SPB:    0
LUN Idle Ticks SPA:    0
LUN Idle Ticks SPB:    0

```

---

---

```
Number of arrivals with non-zero queue:
0
Sum queue lengths by arrivals:
0
Statistics logging start time:
09/28/02 13:08:32.0
Statistics logging current time:
09/28/02 09:40:43.0
Explicit Trespases SPA:
0
Explicit Trespases SPB:
0
Explicit Trespases:                                0
Implicit Trespases SPA:
0
Implicit Trespases SPB:
0
Implicit Trespases:                                0
Non-zero Request Count Arrivals SPA:
0
Non-zero Request Count Arrivals SPB:
0
Non-zero Request Count Arrivals:
0
Sum of Outstanding Requests SPA:      0
Sum of Outstanding Requests SPB:      0
Sum of Outstanding Requests:          0
Bus 0 Enclosure 1 Disk 0 Queue Length:
0
Bus 0 Enclosure 1 Disk 1 Queue Length:
0
Bus 0 Enclosure 1 Disk 2 Queue Length:
0
Bus 0 Enclosure 1 Disk 0 Hard Read
Errors:                                0
Bus 0 Enclosure 1 Disk 1 Hard Read
Errors:                                0
Bus 0 Enclosure 1 Disk 2 Hard Read
Errors:                                0
Bus 0 Enclosure 1 Disk 0 Hard Write
Errors:                                0
Bus 0 Enclosure 1 Disk 1 Hard Write
Errors:                                0
Bus 0 Enclosure 1 Disk 2 Hard Write
Errors:                                0
```

---

---

```
Bus 0 Enclosure 1  Disk 0  Soft Read
Errors:              0
Bus 0 Enclosure 1  Disk 1  Soft Read
Errors:              0
Bus 0 Enclosure 1  Disk 2  Soft Read
Errors:              0
Bus 0 Enclosure 1  Disk 0  Soft Write
Errors:              0
Bus 0 Enclosure 1  Disk 1  Soft Write
Errors:              0
Bus 0 Enclosure 1  Disk 2  Soft Write
Errors:              0

Bus 0 Enclosure 1  Disk 0   Removed
Reads:              0
Writes:              0
Blocks Read:        0
Blocks Written:     0
Queue Max:          N/A
Queue Avg:          N/A
Avg Service Time:  N/A
Prct Idle:          Not Available
Prct Busy:          Not Available
Remapped Sectors:  N/A
Read Retries:      N/A
Write Retries:     N/A

Bus 0 Enclosure 1  Disk 1   Removed
Reads:              0
Writes:              0
Blocks Read:        0
Blocks Written:     0
Queue Max:          N/A
Queue Avg:          N/A
Avg Service Time:  N/A
Prct Idle:          Not Available
Prct Busy:          Not Available
Remapped Sectors:  N/A
Read Retries:      N/A
Write Retries:     N/A

Bus 0 Enclosure 1  Disk 2   Enabled
Reads:              4
Writes:              0
Blocks Read:        28
Blocks Written:     0
Queue Max:          N/A
Queue Avg:          N/A
Avg Service Time:  N/A
Prct Idle:          99.99
Prct Busy:          0.00
```

---

---

```

Remapped Sectors: N/A
Read Retries:      N/A
Write Retries:     N/A
Is Private:        YES
Usage:             Unknown
Snapshots List:    Not Available
MirrorView Name if any: Not
Available

```

```

LOGICAL UNIT NUMBER 4
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.

```

Total Hard Errors:          0
Total Soft Errors:         0
Total Queue Length:        0
Name                        LUN 4
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:          ASAP
Verify Priority:           ASAP

```

---

---

```
Prct Reads Forced Flushed: 0
Prct Writes Forced Flushed: 0
Prct Rebuilt:                100
Prct Bound:                  100
LUN Capacity(Megabytes):    100
LUN Capacity(Blocks):       204800
UID:
60:06:01:60:00:D0:16:00:48:77:A8:54:55:C
2:D6:11
Bus 0 Enclosure 0 Disk 0 Queue Length:
0
Bus 0 Enclosure 0 Disk 1 Queue Length:
0
Bus 0 Enclosure 0 Disk 2 Queue Length:
0
Bus 0 Enclosure 0 Disk 3 Queue Length:
0
Bus 0 Enclosure 0 Disk 4 Queue Length:
0
Bus 0 Enclosure 0 Disk 0 Hard Read
Errors:                0
Bus 0 Enclosure 0 Disk 1 Hard Read
Errors:                0
Bus 0 Enclosure 0 Disk 2 Hard Read
Errors:                0
Bus 0 Enclosure 0 Disk 3 Hard Read
Errors:                0
Bus 0 Enclosure 0 Disk 4 Hard Read
Errors:                0
Bus 0 Enclosure 0 Disk 0 Hard Write
Errors:                0
Bus 0 Enclosure 0 Disk 1 Hard Write
Errors:                0
Bus 0 Enclosure 0 Disk 2 Hard Write
Errors:                0
Bus 0 Enclosure 0 Disk 3 Hard Write
Errors:                0
Bus 0 Enclosure 0 Disk 4 Hard Write
Errors:                0
Bus 0 Enclosure 0 Disk 0 Soft Read
Errors:                0
Bus 0 Enclosure 0 Disk 1 Soft Read
Errors:                0
Bus 0 Enclosure 0 Disk 2 Soft Read
Errors:                0
Bus 0 Enclosure 0 Disk 3 Soft Read
Errors:                0
Bus 0 Enclosure 0 Disk 4 Soft Read
Errors:                0
Bus 0 Enclosure 0 Disk 0 Soft Write
Errors:                0
```

---

---

```
Bus 0 Enclosure 0 Disk 1 Soft Write
Errors:          0
Bus 0 Enclosure 0 Disk 2 Soft Write
Errors:          0
Bus 0 Enclosure 0 Disk 3 Soft Write
Errors:          0
Bus 0 Enclosure 0 Disk 4 Soft Write
Errors:          0

Bus 0 Enclosure 0 Disk 0 Enabled

Bus 0 Enclosure 0 Disk 1 Enabled

Bus 0 Enclosure 0 Disk 2 Enabled

Bus 0 Enclosure 0 Disk 3 Enabled

Bus 0 Enclosure 0 Disk 4 Enabled
Is Private:          NO
Snapshots List:      Not Available
MirrorView Name if any: Not
Available
```

---

**getall -lun command lists only thin LUN information and does not display the private internal LUNs information in user mode.**



```
C:\Program Files\EMC\Navisphere CLI>naviseccli -h 10.14.5.112
-user a -password
a -scope 0 getall -lun
```

```
Server IP Address:      10.14.5.112
Agent Rev:             6.28.10 (1.40)
```

All logical Units Information

-----

```
Statistics logging is disabled.
  Certain fields are not printed if statistics
  logging is not enabled.
LOGICAL UNIT NUMBER 6
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.
```

```
Total Hard Errors:          0
Total Soft Errors:          0
Total Queue Length:         0
Name                         Thin LUN 6
Minimum latency reads N/A
```

```
RAID Type:                  N/A
RAIDGroup ID:               N/A
State:                      Bound
Stripe Crossing:            0
Element Size:               0
Current owner:              SP B
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 B
Rebuild Priority:           ASAP
```

---

```

Verify Priority:                ASAP
Prct Reads Forced Flushed:    0
Prct Writes Forced Flushed:   0
Prct Rebuilt:                  100
Prct Bound:                    100
LUN Capacity(Megabytes):      1024
LUN Capacity(Blocks):         2097152
UID:
60:06:01:60:04:60:1E:00:BE:02:63:43:40:06:DD:11
Bus 0 Enclosure 0 Disk 0 Queue Length:          0
Bus 0 Enclosure 0 Disk 1 Queue Length:          0
Bus 0 Enclosure 0 Disk 2 Queue Length:          0
Bus 0 Enclosure 0 Disk 3 Queue Length:          0
Bus 0 Enclosure 0 Disk 4 Queue Length:          0
Bus 0 Enclosure 0 Disk 0 Hard Read Errors:      0
Bus 0 Enclosure 0 Disk 1 Hard Read Errors:      0
Bus 0 Enclosure 0 Disk 2 Hard Read Errors:      0
Bus 0 Enclosure 0 Disk 3 Hard Read Errors:      0
Bus 0 Enclosure 0 Disk 4 Hard Read Errors:      0
Bus 0 Enclosure 0 Disk 0 Hard Write Errors:     0
Bus 0 Enclosure 0 Disk 1 Hard Write Errors:    0
Bus 0 Enclosure 0 Disk 2 Hard Write Errors:     0
Bus 0 Enclosure 0 Disk 3 Hard Write Errors:     0
Bus 0 Enclosure 0 Disk 4 Hard Write Errors:     0
Bus 0 Enclosure 0 Disk 0 Soft Read Errors:      0
Bus 0 Enclosure 0 Disk 1 Soft Read Errors:      0
Bus 0 Enclosure 0 Disk 2 Soft Read Errors:      0
Bus 0 Enclosure 0 Disk 3 Soft Read Errors:      0

Bus 0 Enclosure 0 Disk 4 Soft Read Errors:      0
Bus 0 Enclosure 0 Disk 0 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 1 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 2 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 3 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 4 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 0 Enabled
Bus 0 Enclosure 0 Disk 1 Enabled
Bus 0 Enclosure 0 Disk 2 Enabled
Bus 0 Enclosure 0 Disk 3 Enabled
Bus 0 Enclosure 0 Disk 4 Enabled
Is Private:                            NO
Snapshots List:                         Not Available
MirrorView Name if any:                 Not Available

```

---

### LUN pool information (same as returned by command reserved -lunpool -list).

```

Name of the SP:                SP A
Total Number of LUNs in Pool:  9
Number of Unallocated LUNs in Pool: 4
Unallocated LUNs:              14, 24, 23, 22
Allocated LUNs:                16, 17, 18, 20, 21
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

```

### RAID group information (same as returned by command getrg).

---

#### All RAID Groups Information

-----

```

RaidGroup ID:      0          RAID group identification number.
RaidGroup Type:    r5         RAID type.
RaidGroup State:   Explicit_Remove Valid_luns Valid States are Invalid,
                                     Expanding, Defragmenting, Halted,
                                     and Busy.

List of disks:
Bus 0 Enclosure 0 Disk 0      Disks used in the RAID group.
  Bus 0 Enclosure 0 Disk 1
  Bus 0 Enclosure 0 Disk 7
  Bus 0 Enclosure 0 Disk 8
  Bus 0 Enclosure 0 Disk 9

List of luns:      0 1 2 3    LUNs currently in RAID group.
Max Number of disks:      16  Max disks allowed in RAID group.
Max Number of luns:      32  Max LUNs allowed in storage
                               system.

Raw Capacity (Blocks):    343391320  Raw capacity in disk blocks.
Logical Capacity (Blocks): 274713056  Logical capacity in blocks.
Free Capacity
(Blocks, non-contiguous): 266324448  Free (unbound and non-contiguous)
capacity of RAID group in blocks.
Free contiguous group of unbound
segments: 266324448           Displays free contiguous group of
unbound segments in blocks.
Defrag/Expand priority:   Medium     May be High, Medium, or Low.
Percent defragmented:     100        Percent complete of defragmenting.
Percent expanded:         100        Displays if LUN expansion enabled
for this RAID group.

```

---

Disk expanding onto:	N/A	Displays disks expanding onto or N/A for not expanding.
Defrag/Expand priority:	Medium	May be High, Medium, or Low.
Percent defragmented:	100	Percent defragmented.
Lun Expansion enabled:	NO	Whether LUN expansion is enabled for this RAID group.
Legal RAID types:	r5	Displays legal RAID types for LUNs to be bound on the RAID group.

**Storage group information (similar to that returned by storagegroup -list).**

Storage Group Name:	Group1	User defined storage group name
Storage Group UID:	28:D0:64:D0:BF:90:D5:11:80:1D:08:00:1B:41:07:5B	32- byte unique identification number.
HBA/SP Pairs:		HBA UID, SP, and SP Port of storage group
HBA UID		SP Name SPPort
-----		-----
20:00:00:00:C9:20:E2:1C:10:00:00:00:C9:20:E2:1C	SP B	1
HLU/ALU Pairs:		Host LUN number (HLU) and actual LUN number (ALU).
HLU Number	ALU Number	
-----	-----	
0	11	
1	12	
2	4	
Shareable:	NO	YES or NO

**SnapView information from -snapviews switch (same as returned by commands snapview -listsnapableluns, -listsessions, and -listsnapshots).**

Snapviews Information		
-----		
*** List of LUNs that can be snapped***		
Name	LUN n	List of LUN numbers for every LUN that can be a source LUN
Name of the session:	sess4_5	Session name
Number of read requests serviced by the Reserved LUN Pool:		Number

---

Total number of read requests on the snapview logical unit:	Number
Number of reads from the TLU:	Number
Number of write requests in the session:	Number
Number of write requests to Reserved LUN Pool that triggered a COW:	Number
Total number of writes requests on the snapview target logical unit:	Number
Number of write requests larger than the chunk size:	Number
List of Target Logical Units: LUN 4 LUN 5	LUN numbers
snap Logical Units Name:	LUN name
snap Logical Units UID:	32-byte unique ID
Session in simulation:	YES, NO, or Not Supported
Session in persistence:	YES or NO
Session creation time:	Date and Time

The following items display for every snapshot, active or inactive

\*\*\* Info of SnapView Snapshots \*\*\*

SnapView logical unit name: snap4	LUN name
SnapView logical unit ID: 60:06:01:60:16:64:08:00:10:54:78:F7:7B:E4:D7:1 1	32-byte unique ID
Target Logical Unit: 4	LUN ID (same for source and snapshot)
State:	Current state; for example, Active

\*\*\* Info of Reserved SnapView Sessions and Snapshots \*\*\*

---

---

Session Name: SANCopy_isc2to12	Session name assigned at session start
Target Logical Unit: 2	LUN ID (same for source and snapshot)
SnapView logical unit name: SANCopy_isc2to12	LUN name

---

**Example**    `naviseccli -h ss1_spa getall`

---

## getarrayuid

### Gets the storage-system unique ID

**Description** The **navisecli getarrayuid** command displays the unique ID (UID) of the storage system that contains the current SP. For a Fibre Channel storage system this is the storage system WWN. For an iSCSI storage system this is also the WWN, but iSCSI also has a target name, the "iqn" identifier, which is a unique identifier. A storage 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 (*sphosts*), then the CLI displays the unique ID of the storage systems containing those SPs. Knowing the UID of an SP's storage system is useful with MirrorView mirroring commands and for other operations.

**Syntax** **getarrayuid** is used with **navisecli** (described on [page 41](#)) as follows:  
**getarrayuid** [-iscsi | -all] [*sphostnames*]

**Prerequisites** For **navisecli**, you must have a user account on the storage system on which you want to execute the command.

**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 storage system on a combo storage system.

**-all**

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

---

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

---

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

**Example** `naviseccli -h ss1_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



## getcache

### Gets cache environment information

**Description** The `naviseccli getcache` command returns information about the SP's caching environment.

An SP's read cache is automatically enabled on powerup if its size is valid. The write cache is automatically enabled on powerup if the cache size is valid and the storage system has the following components:

- ◆ Two SPs
- ◆ A fully charged standby power supply (SPS)
- ◆ Disk modules in 0:0 through 0:8

Caching statistics are meaningful only if the SP's write cache is enabled.

**Syntax** `getcache` is used with `naviseccli` (described on [page 41](#)) as follows:

```
getcache [-asys] [-bsys] [-bbutd] [-bbutt] [-hacv] [-high] [-low]
[-mirror] [-page] [-pdp] [-ppo] [-rm3a] [-rm3b] [-rsta] [-rstb] [-rsza]
[-rszb] [-spa] [-spb] [-state] [-ucp] [-wst] [-wsz]
```

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

**Options** `-asys`

Returns the size of SP A's memory buffer (system buffer space, excludes cache).

`-bsys`

Returns the size of SP B's memory buffer.

`-bbutd`

Returns the day of the week to perform the standby power supply (SPS) self-test.

`-bbutt`

Returns the time of the day to perform the SPS self-test.

**-hacv** (CX3 series, CX series storage systems, and dual-SP AX4-5 series running Navisphere Manager only)

Checks whether the HA cache vault is enabled or disabled. HA cache vault determines the availability of storage-system write caching when a single drive in the cache vault fails. When the HA cache vault is enabled, write caching is disabled if a single vault disk fails. When disabled, write caching is not disabled if a single disk fails.

**-high**

Returns the write cache high watermark. The high watermark is the percentage of dirty pages, which, when reached, causes the SP to begin flushing the cache.

**-low**

Returns the write cache low watermark. The low watermark is the percentage of cache dirty pages that determines when cache flushing stops. When the write cache low watermark is reached during a flush operation, the SP stops flushing the cache.

**-mirror**

Returns the write cache mirrored status. Always enabled for fibre.

**-page**

Returns the size, in KB, of a cache page. This size can be 2, 4, 8, or 16 KB. The default size is 2 KB.

**-pdp**

Returns the percentage of dirty pages currently in cache, that is, pages that have been modified in the SP's write cache, but that have not yet been written to disk. A high percentage of dirty pages means the cache is handling many write requests.

**-ppo**

Returns the percentage of pages owned by the SP.

**-rm3a**

Returns the amount of memory, in MB, reserved for the optimized RAID 3 data buffers on SP A.

**-rm3b**

Returns the amount of memory, in MB, reserved for the optimized RAID 3 data buffers on SP B.

**-rsta**

Returns the current internal state of the read cache on SP A.

**-rstb**

Returns the current internal state of the read cache on SP B.

**-rsza**

Returns, in MB, the amount of read cache memory SP A is allowed to use.

**-rszb**

Returns, in MB, the amount of read cache memory SP B is allowed to use.

**-spa**

Returns the total number of pages in SP A, each page with the cache page size you selected when setting up storage-system caching. This number equals the cache size divided by the cache page size, minus space for checksum tables. If the storage system has two SPs and both are working, they divide the total number of pages between them. If an SP is idle for a long period or fails, the active SP may increase its share of pages.

**-spb**

Returns the total number of pages in SP B, each page with the cache page size you selected when setting up storage-system caching. This number equals the cache size divided by the cache page size, minus space for checksum tables. If the storage system has two SPs and both are working, they divide the total number of pages between them. If an SP is idle for a long period or fails, the active SP may increase its share of pages.

**-state**

Returns the current state of the SP's read or write cache. The possible read cache states are Enabled, Disabling, and Disabled. An SP's read cache is automatically enabled on powerup if the cache size is a valid number and the SP has at least 8 MB of memory. The possible write cache states are Enabled or Disabled, and several transition states, such as Initializing, Enabling, Disabling, Dumping, and Frozen.

**-ucp**

Returns the number of unassigned cache pages (owned by neither A nor B).

**-wst**

Returns the current internal state of the storage system's write cache.

**-wsz**

Returns the write cache size for the storage systems (which are always equal).

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

```

SP Read Cache State           Enabled
SPA Write Cache State         Disabled
SPB Write Cache State         Disabled
Cache Page size:              2
Write Cache Mirrored:         YES
Low Watermark:                60
High Watermark:               90
SPA Cache pages:              0
SPB Cache pages:              0
Unassigned Cache Pages:       0
Read Hit Ratio:                0
Write Hit Ratio:              0
Prct Dirty Cache Pages =      0
Prct Cache Pages Owned =      0
SPA Read Cache State          Enabled
SPB Read Cache State          Enabled
SP Write Cache State          Disabled
SPS Test Day:                  Friday
SPS Test Time:                 14:00
SP A Physical Memory          128
SP B Physical Memory          128
SP A Free Memory               72
SP B Free Memory               72
SPA Read Cache Size =         62
SPB Read Cache Size =         64
SPA Write Cache Size =         0
SPB Write Cache Size =         0

```

```
SPA Optimized Raid 3 Memory Size = 12  
SPB Optimized Raid 3 Memory Size = 12  
HA Cache Vault:                      Not Available
```

**Example** `naviseccli -h ss1_spa getcache`

## getconfig

### Returns storage-system type configuration information

<b>Description</b>	<p>Displays storage-system type configuration information. If you omit switches, it displays all configuration type information. You can set configuration information with <b>setconfig</b>.</p> <p>Use the <b>getsp</b> command to display SP revision and signature information.</p>
<b>Syntax</b>	<p><b>getconfig</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a>) as follows:</p> <pre><b>getconfig</b> [-ptype] [-que] [-rer] [-tn] [-pg8] [-rep]</pre>
<b>Prerequisites</b>	<p>For <b>naviseccli</b>, you must have a user account on the storage system on which you want to execute the command.</p>
<b>Options</b>	<p><b>-ptype</b></p> <p>Displays the storage-system package type.</p> <p><b>-que</b></p> <p>Displays the setting on how queue full status is handled. <b>Yes</b> means that the FLARE software will return device status to the operating system as busy, instead of queue full, when the device queue is full.</p> <p><b>-rer</b></p> <p>Displays the number of recovered errors.</p> <p><b>-tn</b></p> <p>Displays the state of the Target Negotiate bit.</p> <p><b>-pg8</b></p> <p>Displays SCSI mode page 8 information.</p> <p><b>-rep</b></p> <p>Displays periodic error report information: the number of reported errors or N/A.</p>

**Example** `naviseccli -h ss1_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

<b>Description</b>	The <code>naviseccli getcontrol</code> command returns information about the SP performance.
<b>Syntax</b>	<code>getcontrol</code> is used with <code>naviseccli</code> (described on <a href="#">page 41</a> ) as follows: <code>getcontrol [-allib] [-arv] [-busy] [-cbt] [-date] [-day] [-flush] [-idle] [-read] [-rw] [-sc] [-sf] [-sl] [-time] [-write] [-all]</code>
<b>Prerequisites</b>	For <code>naviseccli</code> , you must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<p><b>-allib</b> Returns the idle and busy ticks for all buses.</p> <p><b>-arv</b> Returns the number of times a user request arrived while at least one other request was being processed.</p> <p><b>-busy</b> Returns the percentage of time the SP is busy.</p> <p><b>-cbt</b> Returns the controller idle and busy ticks.</p> <p><b>-date</b> Returns the SP's system date.</p> <p><b>-day</b> Returns the SP's system day.</p> <p><b>-flush</b> Returns the number of times that flushing was turned on by the high watermark, the low watermark, and the Idle unit. This also returns the number of requests to flush the write cache, and the number of write cache blocks flushed.</p> <p><b>-idle</b> Returns the percentage of time the SP is idle.</p>



**-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.

**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:         OFF
SP Read Cache State        Disabled
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: Not Available
Prct Idle: Not Available
System Date: 05/22/2008
Day of the week: Thursday
System Time: 11:31:29
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: 0
Idle_flush_on: 0
Lw_flush_off: 0
Write_cache_flushes: 0
Write_cache_blocks_flushed: 0
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: 0
Controller idle ticks: 0
```

**Example** `naviseccli -h ss1_spa getcontrol`

## getcrus

### Gets CRU state information

**Description** The **naviseccli getcrus** command returns state information on all or selected CRUs (customer replaceable units) in the storage system except for disks (see the **getdisk** command) and selected FRUs/sub FRUs in the storage 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.

**Syntax** **getcrus** is used with **naviseccli** (described on [page 41](#)) as follows:

```
getcrus [-cablingspsa] [-cablingspsb] [-cachecard] [-cpua] [-cpub]
[-dimma] [-dimmb] [-encdrivetype] [-fana] [-fanb] [-fanc] [-fand]
[-fane] [-fanf] [-fang] [-iocarriera] [-iocarrierb] [-ioa] [-iob] [-lcca]
[-lccb] [-lccreva] [-lccrevb] [-lccsna] [-lccsnb] [-mgmta] [-mgmtb]
[-spa] [-spb] [-spsa] [-spsb] [-speeds] [-vsca] [-vsca] [-vscc] [-vscc]
```

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

**Options** **Note:** The command switches **-cpua**, **-cpub**, **-dimma**, **-dimmb**, **-ioa**, **-iob**, **-mgmta**, and **-mgmtb** are supported only on Secure CLI.

#### **-cablingspsa**

Returns the cabling status for SPS A (standby power supply A).

**-cablingspsb**

Returns the cabling status for SPS B (standby power supply B).

**-cachecard** (AX series only)

Returns the hardware state information of the cache memory card.

The following table lists the possible hardware states. See also the **luncache** command.

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 storage system.

**-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** (Not supported on AX4-5 series)

Returns the state information for fan C.

**-fand** (Not supported on AX4-5 series)

Returns the state information for fan D.

**-fane** (Not supported on AX4-5 series)

Returns the state information for fan E.

**-fanf** (Not supported on AX4-5 series)

Returns the state information for fan F.

**-fang** (Not supported on AX4-5 series)

Returns the state information for fan G.

**-iocarriera** (CX4 series only)

Displays information for the I/O carrier device associated with SP A.

**-iocarrierb** (CX4 series only)

Displays information for the I/O carrier device associated with SP B.

**-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 (supported only on Hammerheads)**

Returns the state of management module A.

**-mgmtb (supported only on Hammerheads)**

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.

**-vsca**

Returns the state information on power supply A (A0 for CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems).

**-vscb**

Returns the state information for power supply B (A1 for CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems).

**-vscc** (Not supported on AX4-5 series)

Returns the state information for power supply C (B0 for CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems).

**-vscd** (Not supported on AX4-5 series)

Returns the state information for power supply D (B1 for CX3 model 10 systems, CX3 model 20 systems, CX3 model 40 systems).

**Output**

Sample output follows. Actual output varies depending on the storage-system type and switches used with the **getcrus** command.

```
DPE Enclosure 0
SP A State:                               Present
SP B State:                               Present
Enclosure 0 Fan A State:                   Present
Enclosure 0 Fan B State:                   Present
Enclosure 0 Power A State:                  Present
Enclosure 0 Power B State:                  Present
Enclosure 0 SPS A State:                    Present
Enclosure 0 SPS B State:                    Empty
Enclosure 0 LCC A State:                    Present
Enclosure 0 LCC B State:                    Present

DAE Enclosure 1
Enclosure 1 Fan A State:                     Present
Enclosure 1 Power A State:                   Present
Enclosure 1 Power B State:                   Present
Enclosure 1 LCC A State:                     Present
Enclosure 1 LCC B State:                     Present
```

The **getcrus** display includes the bus number, 0 or 1, for each enclosure.

```
Example  navisecli -h ss1_spa getcrus
           navisecli -h ss1_spa getcrus -iocarriera

SPE2 Enclosure SPE
Enclosure SPE SP A IO Carrier 0 State: Present
Enclosure SPE SP A IO Carrier 0 Power Status: Good
```



## getdisk

### Gets disk status

**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:

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 LUN.
Enabled	Either a hot spare on standby or part of a bound LUN that is assigned to (owned by) the SP you are using as the communication channel to the chassis. If the storage system has another SP, this module's status is Ready when you use the other SP as the communication channel to the chassis.
Equalizing	Data from a hot spare is being copied onto a replacement disk module.
Expanding	Disk is being added to a RAID group.
Failed	Powered down or inaccessible.
Formatting	Being hardware formatted. Generally, modules do not need formatting.
Off	Powered off by the SP, which can happen if a wrong size module is inserted.
Powering Up	Power is being applied to the disk module.
Ready	Module is part of a broken LUN or a LUN that is bound and unassigned. This can mean that the disk module is part of a LUN that is <i>not</i> owned by the SP that you are using as the communication channel to the chassis. If the disk module is part of a LUN assigned to an SP other than the one you are using as the communication channel to the chassis, the module's status is either Enabled or Ready. It is Enabled when you use the other SP as the communication channel to the chassis.
Rebuilding	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or a replacement disk module.

State	Meaning
Removed	Disk module is removed from the chassis; applies only to a disk module that is part of a LUN.
Full Power	Drive is in full-power state; no delay occurs with media access.
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.
Hot Spare Ready	Module is either a hot spare or replacement disk module that replaced a failed module in a LUN. The data is being rebuilt on a hot spare or replacement disk module.
Unbound	Ready to be bound into a LUN.
Unformatted	Disk is unformatted.
Unsupported	Disk is unsupported.

**Syntax** `getdisk` is used with `naviseccli` (described on [page 41](#)) as follows:

```
getdisk [disk-position] [-bind] [-bndcrs] [-busyticks] [-bytrd]
[-bytwrt] [-capacity] [-cpn] [-drivetype] [-hr] [-hs] [-idleticks] [-hw]
[-lun] [-numluns] [-powersavingsdiskcapable]
[-powersavingsdiskslotcapable] [-powersavingsstate] [-private]
[-product] [-rb] [-read] [-rds] [-rev] [-rg] [-sectors] [-serial] [-sig]
[-speeds] [-spinstats] [-sr] [-state] [-sw] [-tla] [-type] [-usercapacity]
[-vendor] [-write] [-wrts] [-all]
```

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

**Options** *disk-position*

Returns status for a specific disk. See “LUN IDs, unique IDs, and disk IDs” on [page 46](#) 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.

**-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.

**-bytrd**

Returns the number of KB read.

**-bytwr**

Returns the number of KB written.

**-capacity**

Returns the disk capacity in MB.

**-cpn**

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

**-drivetype**

Returns the drive type of the disk.

**-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 [page 174](#).

**-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.

**-lun**

Returns the LUN number 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.

**-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.

**-speeds (CX3 series only)**

Displays the current and maximum speed of the disk. The current speed indicates the speed a disk is currently running at; the maximum speed indicates the highest speed a disk is capable of running at. The current speed is dependent on the speed of the enclosure that contains the disk, which is also dependent on the current speed of the bus that the enclosure is running on.

**-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 CLARiiON 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.

**-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, and power saving attributes.

**Output**

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

```
Vendor Id:                SEAGATE
Product Id:               STT14685 CLAR146
Product Revision:        630A
Lun:                      Unbound
Type:                     N/A
State:                    Unbound
Hot Spare:                NO
Prct Rebuilt:             Unbound
Prct Bound:               Unbound
Serial Number:            3LN1WZPD
Sectors:                  0 (0)
Capacity:                 136888
Private:                  Unbound
Bind Signature:           0x0, 0, 4
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:         DG118032553
Request Service Time:         N/A
Read Requests:                0
Write Requests:               0
Kbytes Read:                  0
Kbytes Written:               0
Stripe Boundary Crossing:     None
Drive Type:                   Fibre Channel
Clariion TLA Part Number:     005048730
User Capacity:                0
Idle Ticks:                   0
Busy Ticks:                   0
Current Speed:                4Gbps
Maximum Speed:                4Gbps
Hardware Power Savings       Yes
Capable
Hardware Power Savings Slot   No
capable
Power Savings State           Full Power

```

The following is sample output with the power savings information:

**naviseccli -h 10.14.46.236 getdisk -all**

```

Bus 2 Enclosure 0 Disk 0
Vendor Id:                   SEAGATE
Product Id:                   STT34007 CLAR400
Product Revision:             C003
Lun:                          Unbound
Type:                         N/A
State:                        Unbound
Hot Spare:                    NO
Prct Rebuilt:                 Unbound
Prct Bound:                   Unbound
Serial Number:                3RJ0QQPS
Sectors:                      0 (0)
Capacity:                     375600
Private:                      Unbound
Bind Signature:               0x0, 0, 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:          This disk does not belong to a
RAIDGroup
Clariion Part Number:   DG118032576
Request Service Time:   N/A
Read Requests:          0
Write Requests:         0
Kbytes Read:            0
Kbytes Written:         0
Stripe Boundary Crossing: None
Drive Type:             Fibre Channel
Clariion TLA Part Number:005048775
User Capacity:          0
Idle Ticks:             156837
Busy Ticks:             0
Current Speed: 4Gbps
Maximum Speed: 4Gbps
Queue Max:              N/A
Queue Avg:              N/A
Prct Idle:              1
Prct Busy:              0
Hardware Power Savings Qualified: NO
Hardware Power Savings Eligible: NO
Power Savings State: Full Power
Current Power Savings Log Timestamp: 06/24/09 08:47:38
Spinning Ticks: 0
Standby Ticks: 0
Number of Spin Ups: 0

```

**Example** `naviseccli -h ss1_spa getdisk 1_1_1`

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



## getlog

### Gets an SP or Event Monitor log

#### 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:** Before contacting your service provider, go the CX3 series or CX series self-service website.

#### Syntax

**getlog** is used with **naviseccli** (described on [page 41](#)) as follows:

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

**Note:** For a host not attached (through a switch or HBA) to a storage system, use **getlog -em**.

#### Prerequisites

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

#### 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

**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 storage systems.

---

**Syntax** `getloop` is used with `naviseccli` (described on [page 41](#)) as follows:

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

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

**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 ss1_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

#### Description

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

**Note:** For CX3 series and CX series storage systems, values returned for fields noted with an asterisk (\*) below may not be correct. Use Navisphere Analyzer for the correct values.

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 **metlun -list** command (see [page 343](#)) 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.

#### Syntax

**getlun** is used with **naviseccli** (described on [page 41](#)) as follows:

```
getlun lun-number [-aa] [-addroffset] [-at]
[-bind] [-bp*] [-bread**] [-brwsp *] [-busy**] [-busyticks] [-bwrite**]
[-capacity] [-crus] [-default] [-disk][-drivetype] [-dsa] [-element]
[-ena**] [-etp*] [-ff*] [-hrdscsi] [-idle**] [-idleticks][-idt] [-ismetalun]
[-isthinlun] [-it ] [-itp* ] [-lunbusytickssp*] [-lunidletickssp*]
[-luncache] [-mirrorname] [-mp][-mt ] [-name] [-nanzq*] [-nminus]
[-nzcra*] [-offset] [-owner] [-parityelements] [-pd] [-pic ] [-prb]
[-prefetch] [-prf] [-ps ] [-private] [-psm] [-que] [-rb] [-rc ] [-reads**]
[-ret] [-rg] [-rhist*] [-rwf] [-rwrsp*] [-sc] [-sftscsi] [-slst*] [-slct*]
[-snapshot] [-sor*] [-sqlah*] [-srcp] [-ss] [-ssm] [-state] [-status]
[-stripe] [-totque] [-type] [-uid] [-upb*] [-usage] [-verify]
[-was] [-wc] [-wch*] [-wh*] [-whist*] [-writes**]
```

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

**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 storage system. See [“LUN IDs, unique IDs, and disk IDs” on page 46](#) for information on 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).

**-bind**

Returns the percentage of the disk that is bound.

**-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 storage 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.

**-ff\***

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

**-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.

**-idt**

Returns, in 100-ms units, the idle delay time for caching. For example, 5 equals 500 ms.

**-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.

**-it**

Returns the idle threshold value for caching; that is, the maximum number of I/Os that can be outstanding to a LUN and still have the LUN be considered idle.

**-itp\***

Returns the implicit trespasses in SP A and SP B.

**-lunbusytickssp\***

Returns the amount of time that the disks on this LUN are busy in SP A and SP B.

**-lunidletickssp\***

Returns the amount of time that the disks on this LUN are idle in SP A and SP B.

**-luncache**

Returns LUN offline (cache dirty) condition information for the specified LUN. See also the **luncache** command.

**-mirrorname**

Returns the MirrorView mirror name.

**-mp**

Returns the absolute maximum number of disk blocks to prefetch in response to a read request for variable-length prefetching.



**-name**

Returns the LUN name. This switch applies only to storage systems running a version of FLARE 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 storage systems.

---

**-nanzq\***

Returns the number on arrivals with nonzero queue.

**-nminus**

Returns the status of minimal-latency read support (is or is not supported).

**-nzcra\***

Returns the nonzero request count arrivals in SP A and SP B.

**-offset**

Returns the alignment offset, assigned when the LUN was bound. Refer to the **bind - offset** switch.

**-owner**

Returns the name of the SP that currently owns the LUN. If neither SP owns the LUN a value of **Unknown** is returned.

**-parityelements** (CX3 series only)

Displays the number of parity elements for a LUN with a RAID type of RAID 6 or RAID 5. This switch is relevant only for LUNs that have a RAID type of RAID 6 or RAID 5. For a LUN with a RAID type other than RAID 6 or RAID 5, the CLI displays an N/A value.

**-pd**

Returns, in blocks, the prefetch disable size. The prefetch disable size is the size of the smallest read request for which prefetching is disabled.

**-pic**

Returns the maximum number of I/Os that can be outstanding to the unit and have the unit still perform prefetching.

**-prb**

Returns the percentage of the LUN that is rebuilt.

**-prefetch**

Returns the prefetch type that the LUN is using for read caching. The possible output follows.

Variable length prefetching YES

Constant length prefetching YES

Prefetching: NO

**-prf**

Returns the percentage of read requests that caused a flush.

**-ps**

Returns the number of blocks of data prefetched for one host read request for constant-length prefetching. The prefetch size must be equal to or greater than the segment size. Ranges from 0 to 8192.

**-private**

Returns information on all private LUNs in the storage system.

**-psm**

Returns the variable prefetch size. The variable prefetch size is determined by multiplying the size of the read request, in blocks, by the multiplier. For example, if the prefetch multiplier is set to 4 and the amount of data requested is 2 KB (4 disk blocks), then the variable prefetch size is 4 times 2 KB or 8 KB (16 disk blocks).

**-que**

Returns the queue length, a number.

**-rb**

Returns the value for the rebuild priority (ASAP, High, Medium, or Low).

**-rc**

Returns the LUN's read cache state (enabled or disabled).

**-reads\*\***

Returns the number of reads received for each CRU.

**-ret**

Returns the priority for retaining prefetched data when the read cache becomes full. If the value returned is YES, prefetched data has priority over host-requested data. If the value returned is NO, both data types have the same priority.

**-rg**

Returns the RAID group ID.

**-rhist\***

Returns Read Histogram information. Read Histogram is an array of 10 locations that contain the number of reads. Element  $n$  of the array 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. **-rhist** also displays read histogram overflows, which are the number of I/O operations that were larger than 512 blocks.

**-rwf**

Returns the percentage of write requests that caused a cache flush.

**-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.

**-slst\***

Returns the statistics logging start time.

**-slct\***

Returns the statistics logging current 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.

**-ss**

Returns the prefetch segment size; that is, the size of the segments that make up a constant-length prefetch operation. For constant-length prefetching, **-ss** returns the segment size (in blocks) of data prefetched in one read operation from the LUN.

**-ssm**

Returns the prefetch segment size/multiplier; that is, the size of the segments that make up a prefetch operation. For variable-length prefetching, **-ssm** returns the multiplier which determines the amount of data, relative to the amount of requested data prefetched in one read operation. For example, if the segment multiplier is 4, the segment size is 4 times the amount of data requested.

**-state**

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

**-status** (AX series, CX series, and CX3 series only)

Returns the state of the Device Map for a specified LUN. The Device Map stores the configuration information about all the replication software in the storage system that is associated with the LUN.

**-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

---

**Note:** RAID 6 is available only for CX series storage systems running FLARE 02.26.xxx.5.yyy or later and for CX3 series storage systems running FLARE 03.26.xxx.5.yyy or later.

---

**-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`.

**-verify**

Returns the value for the verify priority (ASAP, High, Medium, or Low).

**-was**

Returns the minimum block size request that bypasses cache.

**-wc**

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

**-wch\***

Returns the number of times that a write was completely satisfied by the cache, avoiding a disk operation.

**-wh\***

Returns the write hit information, if statistics logging is enabled.

**-whist\***

Returns write histogram information. Write Histogram is an array of 10 locations that contain the number of writes. Element *n* of the array 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.

**Example**

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

```
naviseccli -h ss1_spa getlun -addroffset -offset -capacity
```

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

LOGICAL UNIT NUMBER 1
Offset:                0
LUN Capacity(Megabytes): 5120
LUN Capacity(Blocks):  10485760
Address Offset:        2097152

LOGICAL UNIT NUMBER 2
Offset:                7530
LUN Capacity(Megabytes): 10240
LUN Capacity(Blocks):  20971520
Address Offset:        12582912
```

```
naviseccli -h ss1_spa getlun 13
```

This example retrieves information about LUN number 13:

(See sample listing following)

**Output**

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

```
LOGICAL UNIT NUMBER          13
Prefetch size (blocks) =     0
Prefetch multiplier =        48
```

```

Segment size (blocks) =          0
Segment multiplier =            4
Maximum prefetch (blocks) =     512
Prefetch Disable Size (blocks) = \129
Prefetch idle count =           40
Variable length prefetching     YES
Prefetched data retained         YES
Read cache configured according
to specified parameters.

Name                             LUN 13
Minimum Latency Reads            N/A
RAID Type:                       RAID5
RAIDGroup ID:                    0
State:                            Bound
Stripe Crossing:                 0
Element Size:                    128 128
Current owner:                   SP B
Offset:                          0 0
Auto-trespass:                   DISABLED
Auto-assign:                     DISABLED
Write cache:                     ENABLED
Read cache:                      ENABLED
Idle Threshold:                  0 0
Idle Delay Time:                 20
Write Aside Size:                1023
Default Owner:                   SP B
Rebuild Priority:                 ASAP
Verify Priority:                  Low
rct Reads Forced Flushed:        0
Prct Writes Forced Flushed:      0
Prct Rebuilt:                    100
Prct Bound:                      100
LUN Capacity(Megabytes):         1024
LUN Capacity(Blocks):            2097152
UID:                             60:06:01:FD:26:60:00:00:
                                C9:27:93:85:F6:76:D5:11

Enclosure 0 Disk 0 Queue Length: 0
Enclosure 0 Disk 1 Queue Length: 0
Enclosure 0 Disk 7 Queue Length: 0
Enclosure 0 Disk 8 Queue Length: 0
Enclosure 0 Disk 9 Queue Length: 0
Enclosure 0 Disk 1 Hard Read Errors: 0

```

```
Enclosure 0 Disk 7 Hard Read Errors:      0
Enclosure 0 Disk 8 Hard Read Errors:      0
Enclosure 0 Disk 9 Hard Read Errors:      0
Enclosure 0 Disk 0 Hard Write Errors:     0
Enclosure 0 Disk 1 Hard Write Errors:     0
Enclosure 0 Disk 7 Hard Write Errors:     0
Enclosure 0 Disk 8 Hard Write Errors:     0
Enclosure 0 Disk 9 Hard Write Errors:     0
Enclosure 0 Disk 0 Soft Read Errors:      0
Enclosure 0 Disk 1 Soft Read Errors:      0
Enclosure 0 Disk 7 Soft Read Errors:      0
Enclosure 0 Disk 8 Soft Read Errors:      0
Enclosure 0 Disk 9 Soft Read Errors:      0
Enclosure 0 Disk 0 Soft Write Errors:     0
Enclosure 0 Disk 1 Soft Write Errors:     0
Enclosure 0 Disk 7 Soft Write Errors:     0
Enclosure 0 Disk 8 Soft Write Errors:     0
Enclosure 0 Disk 9 Soft Write Errors:     0
Bus 0 Enclosure 0 Disk 0 Enabled
Bus 0 Enclosure 0 Disk 1 Enabled
Bus 0 Enclosure 0 Disk 7 Enabled
Bus 0 Enclosure 0 Disk 8 Enabled
Bus 0 Enclosure 0 Disk 9 Enabled
Is Private:                               NO
Usage:
Mirrored Name if any                      Not Mirrored
```



## getresume

### Displays resume information for storage-system devices

**Description** The **naviseccli getresume** command displays the properties (resume information) for storage-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 specify the device you want to display resume information for, or you can display resume information for all devices.

The product serial number, product part number, and product revision are displayed only for the SPE enclosure on the CX4 series storage system. This information is displayed only when you use the **-all** switch. Only Secure CLI supports use of this switch.

**Syntax** **getresume** is used with **naviseccli** (described on [page 41](#)) as follows:

```
getresume [-all] [-cpu] [-io] [-iocarrier] [-lcc busNumber enclosureNumber [lcca | lccb]] [-mgmt mgmta | mgmtb] [-mp] [-pc] [-ps busNumber enclosureNumber | xpe [psa | psb]] [-sp sp] [-sps]
```

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

### Options

**-all**

Displays resume information for all devices.

**-cpu** (CX3 series only)

Displays resume information for the central processing unit (CPU) modules.

**-io** (CX3 series and AX4-5 series only)

Displays resume information for the personality card.

**-iocarrier** (CX4 series only)

Displays resume information for I/O carrier devices on both SPs.

**-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.

**-mgmt [mgmta | mgmtb]** (CX3 series only)

Displays resume information for the management modules. If you specify **mgmta**, resume information is displayed for the management module associated with SP A; if you specify **mgmtb**, resume information is displayed for the management module associated with SP B.

**-mp**

Displays resume information for the enclosures.

**-pc** (Not supported on AX4-5 series)

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.

---

**Note:** For CX3 model 40 systems, CX3 model 20 systems, and CX3 model 10 systems, resume information is displayed for the power supplies associated with the SP you specify in the command line. If you omit **psa** or **psb** arguments, N/A values are returned for the power supplies that are not associated with the SP you specify.

---

**-sp sp**

Displays resume information for the specified SP.

---

**Note:** For CX3 series and later storage systems, resume information for both SPs is displayed. If the SP specified in the command line is not functional, the CLI continues to return resume information for both SPs.

---

**-sps**

Displays resume information for the standby power supply (SPS) that is associated with the SP you specify in the command line.

**Output**

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

Sample output of **getresume -all** command on a CX4 series storage system for an SPE enclosure.

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: N/A
Programmable Revision: N/A
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: N/A
Programmable Revision: N/A
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: N/A
Programmable Revision: N/A

```

**Example** `naviseccli -h ss1_spa getresume -ps 0 1`

This example displays information stored in the resume PROM for power supplies A and B (bus 0, enclosure 1).

```

Bus 0 Enclosure 1
Power A
EMC Part Number 118032322
EMC Artwork Revision 000
EMC Assembly Revision A03
EMC Serial Number AC115040501930
Vendor Part Number API2SG02-710
Vendor Assembly Number M02
Vendor Serial Number AC1040501930
Vendor Name ACBEL POLYTECH INC.
Location of Manufacture Apex, NC USA
Year of Manufacture 2004
Month of Manufacture 02
Day of Manufacture 13
Assembly Name 12V P/S w/BWR, RPWR
Programmable Name N/A
Programmable Revision N/A
Bus 0 Enclosure 1
Power B
EMC Part Number 118032322
EMC Artwork Revision 000
EMC Assembly Revision A03
EMC Serial Number AC115040501699
Vendor Part Number API2SG02-710
Vendor Assembly Number M02
Vendor Serial Number AC1040501699
Vendor Name ACBEL POLYTECH INC.
Location of Manufacture Apex, NC USA
Year of Manufacture 2004
Month of Manufacture 02
Day of Manufacture 13
Assembly Name 12V P/S w/BWR, RPWR
Programmable Name N/A
Programmable Revision N/A

```

## getrg

### Gets RAID group information

<b>Description</b>	<p>The <b>naviseccli getrg</b> command returns information about the specified RAID group.</p> <p>If no RAID group is specified, the command returns information about all RAID groups.</p>
<b>Syntax</b>	<p><b>getrg</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a>) as follows:</p> <pre><b>getrg</b> <i>rgID</i> [-disks] [-drivetype] [-exdisks] [-hotspare] [-legal] [-lunex] [-lunlist] [-lusc] [-maxd] [-maxl] [-pod] [-powersavingssetting] [-powersavingseligible] [-isrginstandbystate] [-prcntdf] [-prcntex] [-state] [-tcap] [-type] [-ucap] [-all]</pre>
<b>Prerequisites</b>	<p>For <b>naviseccli</b>, you must have a user account on the storage system on which you want to execute the command.</p>
<b>Options</b>	<p><i>rgID</i></p> <p>Is the RAID group identification number. <i>rgID</i> specifies a number between 0 and the maximum number of RAID groups supported by the RAID group storage system.</p> <p><b>-disks</b></p> <p>Returns the state of disks in the RAID group.</p> <p><b>-drivetype</b></p> <p>Returns information about what drive type the RAID group consists of.</p> <p><b>-exdisks</b></p> <p>Returns information about which disks are expanding. Displays N/A for not expanding.</p> <p><b>-hotspare</b></p> <p>Returns the RAID group number that has the faulted disks for which the hot spare is replacing.</p> <p><b>-legal</b></p> <p>Returns information about which RAID types LUNs can be bound as on the specified RAID group.</p>

How the new LUNs can be bound depends on the number of LUNs that already exist and upon the number of disks in the RAID group.

**-lunx**

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. Valid values are high, medium, low, or N/A.

**-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.

**-prcntex**

Returns the percent of expansion that is complete.

**-state**

Returns RAID group state. Valid states are: Invalid, Explicit\_Remove, Valid\_luns, Expanding, Defragmenting, Halted, and Busy.

**-tcap**

Returns the raw and logical capacity of the RAID group in blocks.

**-type**

Returns RAID group type.

**-ucap**

Returns the free (unbounded) capacity of the RAID group in blocks.

**-all**

Lists the RAID group information.

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

```
RAID Group ID:      1
RAID Group Type:   r5
RAID Group State:  Explicit_Remove
                   Valid_luns
List of disks:     Bus 0  Enclosure 0  Disk 0
                   Bus 0  Enclosure 0  Disk 1
                   Bus 0  Enclosure 0  Disk 2
                   Bus 0  Enclosure 0  Disk 3
                   Bus 0  Enclosure 0  Disk 4
                   Bus 0  Enclosure 0  Disk 5
                   Bus 0  Enclosure 1  Disk 0
                   Bus 0  Enclosure 1  Disk 1
                   Bus 0  Enclosure 1  Disk 3
                   Bus 0  Enclosure 1  Disk 2
List of luns:      0 2 3 7
Max Number of disks: 16
Max Number of luns: 32
Raw Capacity (Blocks): 170795880
Logical Capacity (Blocks): 153716292
Free Capacity (Blocks,non-contiguous): 113871492
```



```
Free contiguous group of unbound segments: 113871492
Defrag/Expand priority: Low
Percent defragmented: 100
Percent expanded: 100
Disk expanding onto: N/A
Lun Expansion enabled: NO
Legal RAID types: r5
Power Savings Setting: On
RAID GROUP Power Savings Eligible: NO
Is RAID GROUP in Power Savings Mode: NO
```

**Example** `naviseccli -h ss1_spa getrg 1`

Retrieves information about RAID group 1.

## getsniffer

### Gets background verify reports

**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 storage system.

---

**Note:** The `getsniffer` command does not support thin LUNs.

---

**Syntax** `getsniffer` is used with `naviseccli` (described on [page 41](#)) as follows:

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

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

**Options** `-all`

Retrieves reports on all LUNs in the storage system, owned by the target SP.

---

**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*

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

`-rg raidgroupnumber`

Retrieves reports on all LUNs in the RAID group, owned by the target SP.

`-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.

**-nonv**

Retrieves and displays the report of historical totals of all non-volatile verifies.

**-rec**

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

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

```

VERIFY RESULTS FOR UNIT 1
Sniffing state:ENABLED
Sniffing rate(100 ms/IO):4
Background verify priority:ASAP
Historical Total of all Non-Volatile Recovery Verifies(0
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

Currently Running Full Unit Verify
-----
Verify State: Sniff Running
Percent Complete:0
                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

Most Recently Completed Full Unit Verify
-----
                CorrectedUncorrectable
Checksum errors    0          0
Write Stamp errors 0          0
Time Stamp errors 0          0
Shed Stamp errors 0          0
Coherency errors  0          0

Historical Total of All Full Unit Verifies(0 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

```

**Example** `naviseccli -h ss1_spa getsniffer 1`

Retrieves the report of the active full unit verify operation for LUN 1.

## getsp

### Gets SP revision and serial number information

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

---

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

---

**Syntax** `getsp` is used with `naviseccli` (described on [page 41](#)) as follows:

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

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

**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 SP serial number.

`-mem`

Displays the SP memory size.

`-id`

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

**Example** `naviseccli -h ss1_spa getsp`

SP A

Type of Each SP:	Unknown
Signature For The SP:	147509
Signature For The Peer SP:	135171
Revision Number For The SP:	0.00.00
Serial Number For The SP:	N/A
Memory Size For The SP:	930
SP SCSI ID if Available:	0

SP B

Type of Each SP:	Unknown
Signature For The SP:	135171
Signature For The Peer SP:	147509
Revision Number For The SP:	1.56.20
Serial Number For The SP:	A1000330
Memory Size For The SP:	930
SP SCSI ID if Available:	0

---

## getsptime

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

**Description** The `naviseccli getsptime` command displays the system time set on an SP. Times are the same on SPs, since the FLARE software ensures that both SP clocks show the same time.

**Syntax** `getsptime` is used with `naviseccli` (described on [page 41](#)) as follows:  
`getsptime [-spa | -spb]`

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

**Options** `-spa` or `-spb`  
Specifies from which SP to display the date and time.

**Output** The SP date and time.

**Example** `naviseccli -h data1_spa getsptime -spa`

## getunusedluns

Returns the number of LUNs not used by any storage group or by any replication applications

**Description** The `naviseccli getunusedluns` command displays all the unused LUNs, which are those the storage group and replication applications are not using.

**Syntax** `getunusedluns` is used with `naviseccli` (described on [page 41](#)) as follows:

`getunusedluns`

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

**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
```



## inserttestevent

**Inserts an event into the event monitor log to let you verify the accuracy of an event monitor template**

**Description** The **naviseccli inserttestevent** command inserts a dummy event of code 0x2003 in the Navisphere Event Monitor 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
inserttestevent
```

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

**Example** **naviseccli -h ss1\_spa inserttestevent**

For **ss1\_spa**, this command writes a test event into the Event Monitor log file.

## ioportconfig -list

Displays information about all I/O modules and corresponding I/O ports.

**Description** The `naviseccli ioportconfig -list` command displays the information about all I/O modules and I/O ports on the CX4 series storage system.

**Syntax** `ioportconfig -list` is used with `naviseccli` (described on [page 41](#)) as follows:

`ioportconfig -list -iomodule slot_num [-sp a | b] [-pportid] [-portrole] [-portusage] [-porttype] [-portstate] [-portsubstate] [-isportpersisted] [-lportid] [-all]`

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

### Options

#### **-list**

Displays information about all I/O modules and I/O ports (for both SPs).

#### **-iomodule 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.

#### **-sp a | b**

If the `-sp` switch is not specified, then information about the I/O module with the given `slot_num` for both SPs will be shown.

#### **-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, 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 storage 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.

## Output

---

Information about each I/O module(s) on SPA:

```
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
```

Information about each port on this I/O module:

```
Physical Port ID: 0
Port Role: FE
Logical Port ID: 2
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: 3
Port Usage: Special
Port Type: iSCSI
Port State: Enabled
Port Substate: Good
Is Persisted: Yes
```

---

**Example** `naviseccli -h ss1_spa ioportconfig -list -iomodule 2 -sp a`

This example displays the information for I/O module 2 for SP A.

## ioportconfig -persist

Writes the port configuration information to persistent memory.

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`ioportconfig -persist [-delay delay] [-o]`

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

**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.

**Output** This operation will cause both SPs to reboot in a coordinated fashion. DO YOU WISH TO CONTINUE (Y/N)?

**Example** `naviseccli -h ioportconfig -persist -delay 20`

## **luncache -clear**

**Clears the LUN offline (cache dirty) condition.**

### **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 storage 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 -clear** is used with **naviseccli** (described on [page 41](#)) as follows:

**luncache** *lun-number* **-clear** [-o]

### **Prerequisites**

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

You can verify the offline condition of a LUN, using the **luncache -list** command.

### **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 storage system.

**-o**

Executes the command without prompting for confirmation.

### **Output**

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

**Example** `naviseccli -h ss1_spa luncache 9 -clear`

This command clears the LUN offline condition for the specified LUN.

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)

## **luncache -list**

### **Returns LUN offline (cache dirty) condition information**

**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 storage 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 owns the LUN. You can clear the LUN offline condition using **luncache -clear**. The **luncache -list** command does not support thin LUNs.

---

**Syntax** **luncache -list** is used with **naviseccli** (described on [page 41](#)) as follows:

**luncache *lun-number* -list**

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

**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 storage system.

**Output** LUN Offline (Cache Dirty Condition): YES

**Example** **naviseccli -h ss1\_spa luncache 9 -list**

This command displays LUN offline condition information for the specified LUN.



## managedby

**Displays the application that manages the storage system.**

**Description** The **naviseccli managedby** command displays whether Navisphere Manager or the Navisphere Express application manages the AX4-5 series and AX series storage system you specify.

You use the **managedby** command for purposes of the CLARiiON VSS Provider. The VSS Provider is used in conjunction with the SnapView functionality, to provide backup capabilities. See the *EMC CLARiiON Server Support Products for Windows Installation Guide* for information on the VSS Provider.

If Navisphere Express manages the storage system, the VSS Provider limits some actions. You can use the **managedby** command to determine whether the storage system you specify is managed by Navisphere Manager or Navisphere Express.

**Note:** The command returns an output value of Navi Manager (Navisphere Manager), Navi Express (Navisphere Express), or None. A value of None is synonymous to Navisphere Manager.

**Syntax** **managedby** is used with **naviseccli** (described on [page 41](#)) as follows:

```
managedby
```

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

**Output** Managed By: Navi Manager

**Example** **naviseccli -h ss1\_spa managedby**

This command displays the application that manages the AX series storage system specified. The storage system is managed by Navisphere Manager.

## managefiles -delete

**Deletes SP logging files on the specified storage system.**

**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 to delete using the **-all** or **-file** switches, or you can omit switches and display a list of files, then choose a file from the list.

**Syntax** **managefiles -delete** is used with **naviseccli** (described on [page 41](#)) as follows:

**managefiles -delete** [-all] [-file *filenames*] [-o]

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

**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.

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

**Example** **naviseccli -h ss1\_spa managefiles -delete**

This command shows you all file index numbers and names. The files appear in the format shown below:

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

You then enter the file you want to delete, by either its index number or name.

```
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.
```

## managefiles -list

### Lists SP logging files in supported directories

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`managefiles -list`

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

### 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

**Example** `naviseccli -h ss1_spa managefiles -list`

This command shows you a list of files in the supported directories, for the storage system you specify.

## managefiles -retrieve

Retrieves SP logging files from a remote storage system to the local system

**Description** The **naviseccli managefiles** command with the **-retrieve** function lets you retrieve SP logging files from supported directories on a remote storage system to the local system.

**Note:** The dumps and logs directories are supported for **managefiles**.

You can specify the files to retrieve using the **-all** or **-file** switches, or you can omit switches and display a list of files, then choose a file from the list.

**Syntax** **managefiles -retrieve** is used with **naviseccli** (described on [page 41](#)) as follows:

**managefiles -retrieve [-path *path*] [-all] [-file *filenames*] [-o]**

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

**Options** **-path *path***

Specifies the path on the local system to copy the file to. 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.

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

**Example** `naviseccli -h ss1_spa managefiles -retrieve`

This command shows you all file index numbers and names. The files appear in the format shown below.

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

You then enter the file you want to retrieve, by either its index number or name.

```
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.
```

## ndu -list

### Describes installed SP driver software packages

#### ndu Background

The **ndu** command provides nondisruptive upgrades on CX3 series and CX series storage systems. This command lets you install and upgrade storage-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 FLARE software bundles or enablers.

**Note:** For AX series storage systems the process is disruptive.

Use the **-list** function with no switches to display all the information about all software packages.

**Note:** The **ndu** command is not supported on NetWare and HP-UX operating systems. The **ndu** command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

#### Description

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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
ndu -list [-name name] [-rev] [-iscommitable] [-isrevertable]
[-isactive] [-iscomplete] [-issystem] [-gen]
```

- Prerequisites** For **naviseccli**, you must have a user account on the storage system on which you want to execute the command.
- You must be able to log in to the host running Navisphere CLI.
- 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.
  - iscommittable**  
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.
  - isactive**  
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.
- Output** See previous page. If the version of FLARE software running on the SP does not support this command, a `Not supported` error message is printed to **stderr**.
- Example** **naviseccli -h payroll\_storage\_SPA ndu -list**
- For the SP with hostname **payroll\_storage\_SPA**, this command lists all installed driver packages. For sample output, see previous page.



## ndu -install

### Transfers SP software driver packages to the storage-system private LUN (PSM LUN)

#### Description

The **ndu** command with the **-install** function transfers files from media to the storage system, queries and displays information about the packages, and then installs or upgrades selected software packages.

---

**Note:** For AX series storage systems the process is disruptive. The **ndu** command is not supported on NetWare and HP-UX operating systems. The **ndu** command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

---

The **naviseccli ndu** command **-install** function transfers one or more SP driver packages from a user-accessible file system to the storage-system private storage LUN (PSM). Media should be present before you issue this command.

---

**Note:** Beginning with FLARE Operating Environment version 02.19.xxx.5.yyy, when you execute **ndu -install**, the CLI performs a series of preinstallation validation checks on a CX3 series, CX series, AX4-5 series, or AX series systems, prior to installing or upgrading customer-installable software packages. Prior to issuing the install, you must satisfy the specified configuration requirements for the preinstallation validation checks functionality, or the installation will not complete. See [“Preinstallation validation checks”](#) at the end of this section for configuration requirements.

---

You should use a Windows-based Navisphere client (either Navisphere Manager or Navisphere CLI) 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 storage 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 system that has SnapView, Access Logix, and FLARE 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.

---

### 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:

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 storage system (see [“Getting started with Secure CLI” on page 39](#)). 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 storage system)
    Enter Password
    Enter Scope (0 - global [default]; 1 - local
```

---

**Note:** If the storage 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.

The **ndu** command is not supported on Netware and HP-UX operating systems. The **ndu** command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

---

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.

---

**Syntax** **ndu -install** is used with **naviseccli** (described on [page 41](#)) as follows:

**ndu -install** *pathlist* [-**delay** *seconds*] [-**force**] [-**gen**] [-**verbose**]

**Prerequisites**

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

You must have a user entry in the Navisphere agent configuration file.

**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.

**Output** If the version of FLARE 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
```

**Example** `naviseccli -h ss1_spa ndu -install CX400-02.07.400.3.107.lst`

```
Item number:                                0
Name of the software package: FLARE-Operating-Environment
Revision of the software package:          02.07.400.3.107
Already Installed Revision:                 02.07.400.3.10
Installable                                YES
Disruptive upgrade:                        NO
```

The requested package(s) will be installed. Do you wish to proceed?: (y/n)? **y**

For the SP with hostname **10.14.12.74**, this command installs the FLARE Operating Environment bundle.

## ndu -runrules

### Executes the preinstallation validation checks

#### 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**.

The **ndu** command is not supported on NetWare and HP-UX operating systems. The **ndu** command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

**Note:** Use the **naviseccli ndu -install** command to both run the preinstallation validation checks and install software.

#### Syntax

**ndu -runrules** is used with **naviseccli** (described on [page 41](#)) as follows:

```
ndu -runrules [pathlist] [-listrules] [-verbose]
```

#### Prerequisites

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

You must have a user entry in the Navisphere agent configuration file.

#### 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.

**Output** If the preinstallation validation checks detect failures, the failures are displayed.

**Example** `naviseccli -h ss1_spa ndu -runrules`  
For the SP with hostname **payroll\_storage\_SPA**, this command executes the preinstallation validation checks.

## ndu -commit

**Commits (makes permanent) an installed storage-system driver package**

**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 FLARE and Access Logix software.

The `ndu` command is not supported on NetWare and HP-UX operating systems. The `ndu` command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

Use the `naviseccli ndu -list` command to display installed package information.

---

**Syntax** `ndu -commit` is used with `naviseccli` (described on [page 41](#)) as follows:

`ndu -commit namelist`

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

You must have a user entry in the Navisphere agent configuration file.

**Options** `namelist`

Lists package names of the software packages to be committed.

**Output** If the version of FLARE software running on the SP does not support this command, a `Not supported` error message is printed to `stderr`.

**Example** `naviseccli -h ss1_spa ndu -commit mypackage`

For the SP with hostname `payroll_storage_SPA`, this command commits the installed package `mypackage`.



## ndu -status

Returns the status of the **ndu** install, commit, or revert operations

**Description** The **naviseccli ndu** command **-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.

The **ndu** command is not supported on NetWare and HP-UX operating systems. The **ndu** command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

**Syntax** **ndu -status** is used with **naviseccli** (described on [page 41](#)) as follows:  
**ndu -status [-clear]**

**Prerequisites** For **naviseccli**, you must have a user account on the storage system on which you want to execute the command.  
You must have a user entry in the Navisphere agent configuration file.

**Options** **-clear**  
Clears the status of the last executed **ndu** command.

**Output** If the version of FLARE software running on the SP does not support this command, a `Not supported` error message is printed to **stderr**.

**Example** **naviseccli -h ss1\_spa ndu -status**  
For the SP with hostname **payroll\_storage\_SPA**, this command reports the status of the **ndu -install**, **-commit**, or **-revert** function.

## ndu -revert

**Returns to (restores functionality of) the previous revision of an installed storage-system driver**

### Description

The `naviseccli ndu` command **-revert** function restores the functionality of the previous revision of an installed driver package. A committed package cannot revert, 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.

The `ndu` command is not supported on NetWare and HP-UX operating systems. The `ndu` command is only supported by Secure CLI in Navisphere CLI versions 6.24 and later.

### Syntax

`ndu -revert` is used with `naviseccli` (described on [page 41](#)) as follows:

```
ndu -revert [-delay seconds] namelist [-o]
```

### Prerequisites

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

You must have a user entry in the Navisphere agent configuration file.

### 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.

**Output** If the version of FLARE software running on the SP does not support command, a `Not supported` error message is printed to **stderr**. Other errors:

```
NDU_CANNOT_FIND_FILE
NDU_CANNOT_XFER_FILE
INVALID_PACKAGE
```

**Example** `naviseccli -h ss1_spa ndu -revert mypackage`

For the SP with hostname **payroll\_storage\_SPA**, this command reverts to the previous revision of **mypackage** for both SPs.

## networkadmin -get

### Lists network name and address information

#### 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 storage systems with FLARE 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 **networkadmin -set** function described on [page 251](#).

#### Syntax

**networkadmin -get** is used with **naviseccli** (described on [page 41](#)) as follows:

```
networkadmin -get [-sp a | b] [-name] [-portid portid [-vportid
vportid ] [-vlanid]
[-ipv4 [-mode] [-address] [-subnetmask] [-gateway] ]
[-ipv6 [-mode] [-address] [-globalprefix] [-gateway] ]
[-speed][[-auto] [-capabilities] [-requested] ] [-linkstatus]
[-all]
```

#### Prerequisites

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

- Options**
- sp [a | b]**  
Displays the properties of specified SP. The default is the properties of the connected SP.
  - portid *portid***  
Displays the management port ID. If you do not specify the *portid*, 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 SPs network name.
  - ipv4**  
Displays all the IPv4 information. This is the default IP address.
  - ipv6**  
Displays all the IPv6 information.
  - address**  
Displays the SPs network IP address (IPv4/IPv6 address).
  - gateway**  
Displays the IP address (IPv4/IPv6 address) of the SPs gateway system.
  - subnetmask**  
Displays the SPs IPv4 subnet mask.
  - globalprefix**  
Displays the SPs IPv6 global prefix.

**-mode**

Displays the IPv4/IPv6 status. The statuses are automatic, manual, and disabled.

**-speed** (CX4 series only)

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 does not have a value of the 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 requested speed value.

**-linkstaus** (CX4 series only)

Displays link status of management port.

**-all** (CX4 series only)

Displays all information for **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.

---

**Output**

See the previous page. If the version of FLARE 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**, **-capablespeeds**, **-requestedspeed**, **-linkstatus**, and **-all** are not supported on the storage system then an `Invalid command line parameters error` message is printed to **stderr**.

If IPv6 does not support the storage 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,  
the value will be displayed as an empty string.

**Examples** `naviseccli -h ss1_spa networkadmin -get -speed`

```
Storage Processor SP A
Management Port Settings
Link Status: Link-up
Current Speed: 100 Mbps/half duplex
```

`naviseccli -h ss1_spa networkadmin -get -capabilities -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
```

`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
```

---

**Note:** The IPv6 information contains IPv6 in the header to distinguish the output from the IPv4 information.

---

**naviseccli -h ss1\_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
```



## networkadmin -mib

**Configures the SNMP services on the storage-system SP**

**Description** The **naviseccli networkadmin** command with **-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 -mib** function.

---

**Syntax** **networkadmin -mib** is used with **naviseccli** (described on [page 41](#)) as follows:

**networkadmin -mib [-enable] [-disable] [-community *string*] [-o]**

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

**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.

**Output**

```
Storage Processor:SP A
SNMP MIB Status: Disabled
Community: public
```

If the version of FLARE software running on the SP does not support this command, a Not supported error message is printed to **stderr**.

**Example** `naviseccli -h ss1_spa networkadmin -mib -disable`

Disable processing of SNMP read requests on SP A (y/n) [n]

This command disables the SNMP services on SP A.

## networkadmin -route

Sets up the static routes

**Description** The `naviseccli networkadmin` command with `-route` function is used to set up the static routes to the specific hosts or networks.

**Syntax** `networkadmin -route` is used with `naviseccli` (described on [page 41](#)) as follows:

```
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]
```

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

**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.

**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
```

**Example** `naviseccli -h kirk-spa networkadmin -route -list`

## networkadmin -set

### Updates network information for an SP

#### Description

The **naviseccli networkadmin** command with **-set** function changes one or more network settings for an SP. This command supports IPv6 network communication to the management ports of storage systems with FLARE 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 -get** function.

The network properties are initially set by EMC service personnel 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 use the new value.

#### Syntax

**networkadmin -set** is used with **naviseccli** (described on [page 41](#)) as follows:

```
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]
-gateway gateway -IP-address]]
-speed [-duplex]]
```

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

**Options** **-sp** [*a* | *b*]

Modifies the properties of the specified SP. The default is the connected SP.

**-portid** *portid*

Specifies the management port ID. The default management port is 0.

**-vportid** *vportid*

Specifies the virtual port ID associated with 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 the 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.

**-name** *name*

Changes the SPs network name to *name*. The maximum size of the name is 64 characters. This parameter change restarts the CIMOM.

**-address** *IP-address*

Changes the SPs IPv4 network address to *IP address*.

**-ipv4**

Specifies the settings for IPv4.

**-ipv6**

Specifies the settings for IPv6.

**-automatic**

Enables IPv6 on the SP automatically.

**-disable**

Disables IPv6 on the SP.

**-manual**

Enables IPv6 on the SP. The global prefix and gateway must be set manually.

**-gateway** *gateway-IP-address*

Changes the SPs IPv4 or IPv6 gateway IP address to *gateway-IP-address*.

**-subnetmask** *mask*

Changes the SPs IPv4 subnet mask to *mask*.

**-globalprefix** *prefix*

Changes the SPs IPv6 global prefix to *prefix*.

**-speed [-duplex]** (CX4 series only)

Changes the speed on the management port for the target SP.

**-o**

Executes the command without prompting for confirmation.

**Output**

See the description section of this command. If the version of FLARE 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
```

**Example**

```
naviseccli -h ss1_spa networkadmin -set -gatewayaddress
```

Changing the gateway address of SP A from 123.456.789.011 to 123.456.789.011 (y/n) [n] **y**

SP A reboots.

For the SP with hostname **payroll\_storage\_SPA** this command changes the gateway address.



```
naviseccli -address 10.5.2.12 -user a -password a --scope 0  
networkadmin --set --ipv6 -manual
```

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 -globalprefix 3ffe:80c0:22c:4c:
```

This command enables IPv6 (if not enabled) and sets the IPv6 global prefix.

```
naviseccli -h ss1_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)?
```

For SP A, this command changes the management port speed to 1000 Mbps/half duplex.

## port

**Lists SP port information, removes an HBA entry from an SP's initiator list, and provides diagnostic capabilities**

**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 starting on [page 315](#) and `alpa` command on [page 57](#).

To register an HBA with a storage system (opposite of removing an HBA entry), use the `register` command.

**Syntax** `port` is used with `naviseccli` (described on [page 41](#)) as follows:

```
port [-diagnose [optional-diagnose-switches]] [-list
[optional-list-switches]] [-removeHBA [optional-removeHBA -switches]
[-o]
-list [ [-all] [-arraycommpath] [-arrayhba] [-bread] [-failovermode]
[-gname] [-hba] [-initiatorcount] [-mac] [-reads] [-sfpstate] [-sp]
[-uid StorageGroupUID] [-unitserialnumber] [-writes] [ -physical]
[-usage] ]
-removeHBA [-hbauid hbauid | -host hostname | -hostip IPAdress
|-all]
```

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

**Options** The optional `port -diagnose` switches are:

```
[-sancopy [-clean]] [-host]
-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.

The optional **port -list** switches are:

Without arguments, **-list** displays information about all ports in the system but does not list information given by **-all**, **-arraycomppath**, and **-failovermode**. This allows **-list** without an argument to produce output compatible with scripts written for Navisphere revisions 5.X. Example 1 shows output without switches. Example 2 shows CLI 6.0.5 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 Navisphere release. For script compatibility with revisions of Navisphere before 5.3, use the **-list** without an argument.

**-arraycomppath**

Displays the status of the **arraycomppath** setting: 0 if **arraycomppath** is disabled; 1 if **arraycomppath** is enabled.

**-arrayhba**

Displays information about all storage systems attached to this HBA, or about storage 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** *StorageGroupName*

Specifies a storage group. *StorageGroupName* 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 an iSCSI HBA. When you target a Fibre Channel 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 storage system. 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**(CX3 series only)

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** *StorageGroupUID*

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.

**-unitserialnumber**

Specifies the **unitserialnumber** mode.

**-writes**

Displays the number of writes for each port.

The optional **port -removeHBA** switches are:

**-hbauid** *hbauid*

Removes a specific initiator record specified with **-hbauid**. *hbauid* specifies a unique host bus adapter identification number (must use uppercase).

**-host** *hostname*

Removes all initiator records that are logged out and registered, for the specified hostname.

**-hostip** *IPAddress*

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 storage 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:

**naviseccli -h ss1\_spa port -list**

Information about each HBA:

```
HBA UID:
  10:10:10:10:10:10:10:10:10:10:10:10:10:10:10:10
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
StorageGroup Name:     None
  
```

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 B
SP Port ID:         0
SP UID:
    50:06:01:60:D0:C3:0F:E0:50:06:01:68:50:C3:0F:E0
Link Status:        Up
Port Status:        Online
Switch Present:     YES
Switch UID:
    10:00:00:60:69:10:18:16:20:09:00:60:69:50:14:13
SP Source ID:       727296
```

```
SP Name:            SP B
SP Port ID:         1
SP UID:
    50:06:01:60:D0:C3:0F:E0:50:06:01:69:50:C3:0F:E0
Link Status:        Down
Port Status:        DISABLED
Switch Present:     YES
Switch UID:
    00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
SP Source ID:       239
```

```
SP Name:            SP A
SP Port ID:         0
SP UID:
    50:06:01:60:D0:C3:0F:E0:50:06:01:60:50:C3:0F:E0
Link Status:        Up
Port Status:        Online
Switch Present:     YES
Switch UID:
    10:00:00:60:69:50:10:4A:20:09:00:60:69:50:15:23
SP Source ID:       661760
```

```
SP Name:            SP A
SP Port ID:         1
```



```

SP UID:
    50:06:01:60:D0:C3:0F:E0:50:06:01:61:50:C3:0F:E0
Link Status:          Down
Port Status:          DISABLED
Switch Present:       YES
Switch UID:
    00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
SP Source ID:         239

```

**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.

**Example 2** The following example shows sample output of the `port -list -all` command targeted to Fibre Channel HBAs and SPs:

```
naviseccli -h ss1_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* storage 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* storage 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 -all` command targeted to a CX4 series storage system:

```
naviseccli -h ss1_spa port -list -sp -all
```

```
Information about each SPPORT:
```

```
SP Name: SP A
SP Port ID: 1
SP UID: 50:06:01:60:B0:60:0B:2D:50:06:01:61:30:60:0B:2D
Link Status: Down
Port Status: DISABLED
Switch Present: NO
MAC Address: Not Applicable
SFP State: N/A
I/O Module Slot: 1
Physical Port ID: 0
Usage: MirrorView
```

```
SP Name: SP A
SP Port ID: 0
SP UID: 50:06:01:60:B0:60:0B:2D:50:06:01:60:30:60:0B:2D
Link Status: Up
Port Status: Online
Switch Present: YES
Switch UID:
10:00:08:00:88:E3:20:D8:20:08:08:00:88:E3:20:D8
SP Source ID: 6424595
MAC Address: Not Applicable
SFP State: N/A
I/O Module Slot: 1
Physical Port ID: 1
Usage: General
```

## powersaving

**Enables or disables the power saving settings**

- Description** The `naviseccli powersaving` command allows you to enable or disable power saving settings on a storage system. It displays the system's idle time and the status of the power saving settings.
- Syntax** `powersaving` is used with `naviseccli` (described on [page 41](#)) as follows:
- ```
naviseccli powersaving
-globalsettings [-o] [on | off]
-info [-systemidletime] [-globalpowersavingsettings]
```
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options** `-globalsettings`
- Enables or disables power saving settings on a storage 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

**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 Navisphere agent process may hang. If the agent process hangs, you must reboot the host on which the Navisphere 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. If write cache is enabled, rebooting of both SPs is not allowed on CX4 series storage systems.

---



---

**Note:** See the **rebootpeerSP** command, to reboot the peer SP of the target SP.

---

**Syntax** **rebootSP** is used with **naviseccli** (described on [page 41](#)) as follows:

**rebootSP** [-both] [-o]

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

**Options** -both

Reboot both SPs.

-o

Executes the command without prompting for confirmation.

**Example** `naviseccli -h ss1_spa rebootSP -both -o`



## rebootpeerSP

### Reboots a peer SP

**Description** Lets you reboot the peer SP of the target SP in a dual-SP storage system. For the **rebootpeerSP** command, the peer SP is the SP on the storage system to which the command is not targeted.

---

**Note:** See the **rebootSP** command, to reboot the target SP or both SPs on the storage system.

---

**Syntax** **rebootpeerSP** is used with **naviseccli** (described on [page 41](#)) as follows:

**rebootpeerSP [-o]**

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

**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
```

```
This operation will cause a storage system reboot!
```

```
DO YOU WISH TO CONTINUE? (y/n)?
```

## remoteconfig -getconfig

Displays information about a host or SP agent

### remoteconfig Background

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 storage 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 Navisphere Host Agent is supported only on ESX Server.

### Description

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 528](#). 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
remoteconfig -getconfig [-description] [-contact] [-users] [-interval] [-write filename]
```

### Prerequisites

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

### Options

**-description**

Displays contents of description field.

**-contact**

Displays contents of contact field.

**-users**

Displays contents of users field.

**-interval**

Displays contents of 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**.

**Output** See [remoteconfig -getconfig 271](#). If the version of FLARE 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
```

**Example** `naviseccli -h ss1_spa remoteconfig -getconfig`

This command lists all agent information.

## remoteconfig -setconfig

### Changes configuration information for a host or SP agent

**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 528.

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
remoteconfig -setconfig [-o] [-adduser userlist] [-contact contactinfo]
[-description description] [-f filename][-interval interval] [-rmuser
userlist]
```

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

**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* (Not supported on AX series)

Reads the configuration information from file *filename*. This file must be in the correct format (same as the agent configuration file used in earlier Navisphere 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 storage-system types.

**-rmuser** *userlist*

Removes the specified users from the authorized user list.

**Output**

See [remoteconfig -getconfig 271](#). If the version of FLARE 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
```

**Example**

```
naviseccli -h ss1_spa remoteconfig -setconfig -adduser
tsjones@viking
```

```
Add tsjones@viking in the authorized user list <y/n>[y]?
```

This command adds the user **tsjones** on host **viking** to the agent privileged users list.

---

## removerg

### Removes a RAID group

**Description** The `naviseccli removerg` command removes a specified RAID group.

---

**Note:** This command cannot remove thin pools.

---

**Syntax** `removerg` is used with `naviseccli` (described on [page 41](#)) as follows:

`removerg rgID`

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

**Options** *rgID*

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

**Example** `naviseccli -h ss1_spa removerg 1`

This command removes RAID group 1.

## resetandhold

**Resets an SP and holds the reboot following the power on self tests (POST)**

**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 storage 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** is used with **naviseccli** (described on [page 41](#)) as follows:

**resetandhold**

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

**Example** The following example resets and holds the reboot for the targeted SP on the storage system.

```
naviseccli -h ss1_spa resetandhold
```

```
This operation will cause the storage 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

#### Description

The **naviseccli responstest** command creates an artificial event to test Navisphere 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 storage systems only. The EMC phone home feature is not available via this command.

#### Syntax

**responsetest** is used with **naviseccli** (described on [page 41](#)) with one of the following arguments as follows:

**responsetest -email** *response-parameters* |

**-emailpage** *response-parameters* | **-modempage** *response-parameters* |

**-snmp** *response-parameters*

#### Prerequisites

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

#### 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* ]  
Specifies 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* is the PC COM port number. 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 used by the agent. The **initcommand** allows for some custom modem configuration before the data is sent to it. Optional.
- snmp**
- 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.

**Output** Result of the test.

**Example** `naviseccli -h ss1_spa responsetest modempage -destnumber -msgnumber`

This command sends a modem page test event.

## sc\_off

### Disables data access control

**Description** The `naviseccli sc_off` command disables data access control and turns off shared storage features. You can use this command to reset the storage system to its factory unshared storage settings.



### CAUTION

**Be aware that executing `sc_off` resets the storage system to its factory unshared storage settings, permanently destroying all previously implemented shared storage features.**

**Syntax** `sc_off` is used with `naviseccli` (described on [page 41](#)) as follows:  
`sc_off [-o]`

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

**Options** `-o`  
Executes the command without prompting for confirmation.

**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 storage system.
```

```
The following functions will be performed, if necessary:
```

```
Clear the storage 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 storage 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 storage system.
```

```
The following functions will be performed, if necessary:
```

```
Clear the storage 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
```

**Example** `naviseccli -h ss1_spa sc_off`

## setcache

### Sets SP cache or RAID 3 size information

#### Description

The `naviseccli setcache` command changes system caching or RAID 3 settings on an SP. By using `setcache` with its various switches, you can enable, disable, and configure the cache. Cache size switches do not affect a cache if that cache is enabled; you must disable the pertinent cache (`-rc` or `-wc` switches) before any size switches will be effective.

---

**Note:** The changes to caching properties may be delayed until you have allocated memory and the standby power supply (SPS) is working. This may take several minutes.

---



---

**Note:** You can use the `getsp` command to discover SP memory and the `getcache` command to discover cache settings.

---

#### Syntax

`setcache` is used with `naviseccli` (described on [page 41](#)) as follows:

```
setcache [-h high-watermark] [-hacv n] [-l low-watermark]
[-p page-size] [-r3a raid3-size-spa] [-r3b raid3-size-spb]
[-rsza read-cache-size-spa] [-rszb read-cache-size-spb]
[-rca read-cache-spa] [-rcb read-cache-spb]
[-wc write-cache] [-wsz write-cache-size] [-o]
```

#### Prerequisites

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

#### Options

`-h` *high-watermark* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Sets the write cache high watermark. *high-watermark* specifies the percentage of dirty pages, which, when reached, causes both SPs to begin flushing the cache. The default is 80 percent for CX3 series and CX series.

A lower value causes the SPs to start flushing the cache sooner. The high watermark cannot be less than the low watermark. To turn watermark processing off, set both the low and high watermark values to 100.

**-hacv** *n* (not supported on a single-SP AX4-5 series system running Navisphere Manager and on CX4 series systems)

Enables or disables the HA cache vault. *n* is as follows:

**1** = Enables HA cache vault (the default)

**0** = Disables HA cache vault

If you set **-hacv** to a value other than 0 or 1, an error message appears. If you enable the HA cache vault, a single drive failure will cause the write cache to become disabled, thus reducing the risk of losing data in the event of a second drive failing. If you disable the HA cache vault, a single drive failure does not disable the write cache, leaving data at risk if a second drive fails. When you disable the HA cache vault, you will receive a warning message stating that this operation will allow write caching to continue even if one of the cache vault drives fails. If there is already a failure on one of the cache vault drives, this operation will not re-enable the write cache.

**-l** *low-watermark* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Sets the write cache low watermark. *low-watermark* specifies the percentage of cache dirty pages that determines when cache flushing stops. When the low watermark is reached during a flush operation, both SPs stop flushing the cache. The default is 60 percent for CX3 series and CX series. The low watermark cannot be greater than the high watermark. To turn watermark processing off, set both the low and high watermark values to 100.

**-p** *page-size*

Sets, in KB, the page size for the caches. Valid sizes are 2, 4, 8, and 16. The default is 2. Page size is set to 2 regardless of enable/disable. As a general guideline, the cache page size should be 8 KB for general UNIX file server applications, and 2 KB or 4 KB for database applications. It does not work if the read or write cache is enabled.

---

**Note:** Setting a page size, RAID 3 size, or write cache size for one SP automatically sets the other to the same size, so you need only specify a size for one SP.

---

**-r3a** *raid3-size-spa* (not supported on an AX4-5 series system running Navisphere Manager)

For a non-4700 storage system: Sets the SP memory size, in MB, reserved for RAID 3 in SP A (and SP B). *raid3-size-spa* specifies the size. This size must be the sum of the memory allocated for all RAID 3 LUNs to be bound. For example, if you bind two RAID 3 LUNs, each with the recommended 6 MB of memory, you must set this field to 12 (6 + 6). If the storage system does not have any RAID 3 LUNs, this value should be 0.

For a CX3 series or CX series, RAID 3 memory is allocated automatically; you cannot set it.

**-r3b** *raid3-size-spb* (not supported on an AX4-5 series system running Navisphere Manager)

Sets the optimum RAID 3 size for SP B (and SP A). See comments for **-r3a** above.

**-rsza** *read-cache-size-spa*

Sets the read cache size for SP A to *read-cache-size-spa* MB.

---

**Note:** For read/write caching to function, the total size allocated to the read and write caches (together) must be at least 4 MB (for CX3 series and CX series storage systems).

---

**-rszb** *read-cache-size-spb* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Sets the read cache size for SP B to *read-cache-size-spb* MB.

**-rca** *read-cache-spa*

Enables or disables the read cache function for SP A. For *read-cache-spa*:

- 0 = Disables read cache
- 1 = Enables read cache

**-rcb** *read-cache-spb* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Enables or disables the read cache function for SP B. For *read-cache-spb*:

- 0 = Disables read cache
- 1 = Enables read cache



**-wc** *write-cache* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Enables or disables the write cache function. For *write-cache*:

**0** = Disables write cache

**1** = Enables write cache

**-wsz** *write-cache-size-sp* (not supported on a single-SP AX4-5 series system running Navisphere Manager)

Sets the write cache size for the storage system to *write-cache-size-sp* MB. Also, see the note on sizes under **-rsza**.

### Examples

```
naviseccli -h ss1_spa getsp -mem
```

```
Memory Size For The SP:          930
```

The **getsp -mem** command gets SP memory information. Each SP has the same amount of memory, so you need not issue this command for SP B.

```
naviseccli -h ss1_spa setcache -wsz 650 -wc 1 -p 8
-rcza 280 -rczb 280 -rc 1
```

This command sets the write cache (applies to both SPs) to 650 MB, enables write caching, sets a write cache page size of 8 KB, sets a read cache size for each SP of 280 MB, and enables read caching.

## setfeature

**Adds a software feature to a LUN controlled by an SP**

**Description** In Navisphere, the **naviseccli setfeature** command adds and removes a software driver feature to a LUN owned by an SP.

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

**Syntax** **setfeature** is used with **naviseccli** (described on [page 41](#)) as follows:

**setfeature -on|-off -feature *feature* -lun *lun*|-lunuid *uid***

where:

**-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 FLARE 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**

**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 storage system.

**Note:** The `setsniffer` command does not support thin LUNs.

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

**Syntax** `setsniffer` is used with `naviseccli` (described on [page 41](#)) as follows:  
`setsniffer -all | lun | -rg raidgroupnumber [-bv] [-bvtime priority] [-cr]`

where:

**-all**

Indicates to apply sniffer parameters to all luns in the storage system. The target SP must own one LUN at minimum.

*lun*

Indicates to apply sniffer parameters to the specified LUN. *lun* specifies the logical unit number. You must target the SP that owns the specified LUN.

**-rg** *raidgroupnumber*

Indicates to apply sniffer parameters to all LUNs in the specified RAID group.

**-bv**

Starts a full unit background verify.

**-bvertime** *priority*

Specifies the background verify priority. *priority* specifies the priority value (the rate at which the full background verify is executed). Valid values are ASAP, High, Medium, and Low.

**-cr**

Clears background verify reports.

**Output** There is no output. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa setsniffer 0 -bv -bvertime high`

Starts a background verify on LUN 0 and sets the verify priority to high.

## setspstime

### Sets SPS or BBU battery test time

**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.

**Note:** For AX4-5 series systems, the battery test time on the peer SPS is set to the time you are setting for the current SPS, with an offset of approximately 20 minutes.

**Syntax** `setspstime` is used with `naviseccli` (described on [page 41](#)) as follows:  
`setspstime [-d day-num] [-h hour-num] [-m minute-num] [-nolocal]`

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

**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.

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

**Example** `naviseccli -h ss1_spa setspstime -d 0 -h 1 -m 0`  
 This example sets the SPS battery test time to Sunday at 1:00 AM.

## setsptime

Sets the SP system clocks

**Description** The `naviseccli setsptime` command sets the system time for both SPs.

**Syntax** `setsptime` is used with `naviseccli` (described on [page 41](#)) as follows:  
`setsptime [-M month-num] [-D day-num] [-Y year-num]  
[-h hour-num] [-m minute-num]`

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

**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).

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

**Example** `naviseccli -h data1_spa setsptime -M 3 -D 13 -Y 2002 -h 9`  
This example sets the SP system clock to March 13, 2002, 9:00 am.

## setstats

### Sets statistics logging

**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 storage-system caching that you can turn on and off. When enabled, logging affects storage-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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli setstats -on | -off
[-powersavings] [-on | -off]
```

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

**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.

**Output** There is no output. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa setstats -on`  
`naviseccli -h ss1_spa setstats -powersavings`

The sample output is as follows:

```
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
```



## shutdown

### Shuts down and powers off the storage system

**Description** The **naviseccli shutdown** command shuts down and powers off the storage system you specify.

**Syntax** **shutdown** is used with **naviseccli** (described on [page 41](#)) as follows:  
**shutdown [-o]**

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

**Options** **-o**  
Executes the command without prompting for confirmation.

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

**Example** **naviseccli -h ss1\_spa shutdown**

This command shuts down and powers off the specified storage system.

Before you shutdown the storage system, stop all I/O from any servers connected to the storage 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 storage system you will regain access to the LUNs.

Do you want to shutdown and power off the subsystem now (y/n)?

## shutdownsp

### Shuts down and powers off the current SP

<b>Description</b>	The <b>naviseccli shutdownsp</b> command shuts down the current storage processor (SP) of the storage system.
<b>Syntax</b>	<b>shutdownsp</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a> ) as follows:  <b>shutdownsp [-o]</b>
<b>Prerequisites</b>	For <b>naviseccli</b> , you must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<b>-o</b>  Executes the command without prompting for confirmation. The confirmation message varies depending on the storage system.
<b>Output</b>	None if the command succeeds; error information if used with any other storage system.
<b>Example</b>	<b>naviseccli -h ss1_spa shutdownsp</b>  This command shuts down and powers off the current storage processor (SP).  The confirmation message for AX4-5 series storage system is as follows:  This operation will shutdown and power off the SP.  To reconnect please reboot the SP or manually power up by pushing the power button on the system. Before you shut down the SP, first verify that each server has failover software and has a working path to a port on both SPA and SPB. If you do not, stop I/O from the server(s) without a path to each SP, to prevent losing access to your data.  Do you want to shutdown and power off the SP now? (y/n)?  The confirmation message for CX4 series storage system is as follows:  This operation will power off the SP. To reconnect please reboot the SP - this can be accomplished by using the <b>rebootpeerSP</b> command.  Do you want to shutdown and power off the SP now? (y/n)?

## shutdownpeersp

### Shuts down and powers off the peer SP

- Description** The `naviseccli shutdownpeersp` command shuts down the peer storage processor (SP).
- Syntax** `shutdownpeersp` is used with `naviseccli` (described on [page 41](#)) as follows:  
`shutdownpeersp [-o]`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options** `-o`  
 Executes the command without prompting for confirmation. The confirmation message varies depending on the storage system.
- Output** None if the command succeeds; error information if used with any other storage system.
- Example** `naviseccli -h ss1_spa shutdownpeersp`  
 This command shuts down and powers off the specified peer SP.  
 The confirmation message for AX4-5 series storage system is as follows:
- ```
This operation will shutdown and power off the SP.
To reconnect please reboot the SP or manually power up by
pushing the power button on the system. Before you shut
down the SP, first verify that each server has failover
software and has a working path to a port on both SPA and
SPB. If you do not, stop I/O from the server(s) without
a path to each SP, to prevent losing access to your data.
Do you want to shutdown and power off the SP now? (y/n)?
```
- The confirmation message for CX4 series storage system is as follows:
- ```
This operation will power off the SP.
To reconnect please reboot the SP - this can be
accomplished by using the rebootpeerSP command.
Do you want to shutdown and power off the SP now? (y/n)?
```

## sppcollect -info

### Displays the setting of the automated SPCollect service

**Description** The `naviseccli sppcollect` 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 FLARE Operating Environment version 02.19.xxx.5.yyy or higher.

**Syntax** `sppcollect -info` is used with `naviseccli` (described on [page 41](#)) as follows:

`sppcollect -info`

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

**Output** `AutoExecution: Enabled`

**Example** `naviseccli -address ss1_spa sppcollect -info`

This command shows you the setting of the automated SPCollect functionality for the specified storage system. The *Enabled* setting, indicating that the SPCollect process is turned on, appears in the output above.

## sppcollect -set

Sets the automated SPCollect service on or off

**Description** The `naviseccli sppcollect` 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 FLARE Operating Environment version 02.19.xxx.5.yyy or higher.

---

**Note:** Use `sppcollect -info` to verify whether the service is enabled or disabled.

**Syntax** `sppcollect -set` is used with `naviseccli` (described on [page 41](#)) as follows:

`sppcollect -set -auto on | off [-o]`

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

**Options** `-auto on | off`  
Specifies to set the automated SPCollect process on or off. If you indicate *on*, it turns the automation service on; *off* indicates to turn the automation service off.

`-o`

Executes the command without prompting for confirmation.

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

**Example** `naviseccli -address ss1_spa sppcollect -set -auto on`

This command turns the automated SPCollect process on.

---

## sppcollect

Selects and places a collection of storage system log files into a single .zip file

**Description** The **naviseccli sppcollect** command selects a collection of storage system log files and places them in a single .zip file on the storage system. You can retrieve the file from the storage 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 FLARE Operating Environment 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.

---

**Important:** The SPCollect functionality can affect system performance (may degrade storage system performance).

**Syntax** **sppcollect** is used with **naviseccli** (described on [page 41](#)) as follows:  
**sppcollect**

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

**Output** There is no output. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa spcollect`

This example creates a .zip file on the storage system, which contains system log files.

The .zip file has the following format:

*arrayserialnumber\_SPx\_date\_time\_spsignature\_data.zip*

where:

<i>arrayserialnumber</i>	Is the storage system serial number.
<i>x</i>	Is SP A or B.
<i>date</i>	Is the date of collection.
<i>time</i>	Is the time of collection.
<i>spsignature</i>	Is the SP signature

## sportspeed

### Set or get the SP port speed

#### 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 storage-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 storage system,
supports the new speed. If the speed settings are
incompatible, the storage 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
sportspeed -get [-isauto] [-avail] [-requested] [-type] [-physical]
```

#### Prerequisites

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

#### 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 **sportspeed -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** (CX4 series only)

Displays the physical location of the port.

or

**sportspeed -set -sp** *sp* **-portid** *port speed*

---

**Note:** If an incompatible small form factor pluggable (SFP) is detected, the port is faulted and you will not be able to set the SP port speed.

---

where:

**-sp** *sp*

Specifies the SP to set the port speed on **a** | **b**.

**-portid** *port speed*

*port* 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.

**Output** If the version of FLARE software running on the SP does not support this command, a `Not supported` error message is printed to **stderr**.

**Example** `naviseccli -h ss1_spa sportspeed -get`

```
Storage Processor : SP A
Port ID : 0
Speed Value : 1
```

```
Storage Processor : SP A
Port ID : 1
Speed Value : 2
```

```
Storage Processor : SP B
Port ID : 0
Speed Value : 2
```

```
Storage Processor : SP B
Port ID : 1
Speed Value : 2
```

`naviseccli -h ss1_spa sportspeed -get -physical`

```
Storage Processor :SP A
Port ID :0
Speed Value :4Gbps
I/O Module Slot :1
Physical Port ID :3
```

```
Storage Processor :SP A
Port ID :1
Speed Value :N/A
I/O Module Slot :0
Physical Port ID :2
```

```
Storage Processor :SP B
Port ID :0
Speed Value :4Gbps
I/O Module Slot :1
Physical Port ID :1
```

```
Storage Processor :SP B
Port ID :1
Speed Value :N/A
I/O Module Slot :0
Physical Port ID :3
```

```
naviseccli -h ss1_spa sportspeed -set -sp a -portid 0 2
```

Changing the port speed is a disruptive function. Verify that the other device connected to the storage system, supports the new speed. If the speed settings are incompatible, the storage 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
```

The **get** command example displays the port speeds of the SP **ss1\_spa** and its peer SP. The **set** command example sets the SP port speed of SP A port 0 to 2 Gb/s.

## storagegroup

### Creates and configures storage groups

#### Description

The **naviseccli storagegroup** command lets you create and manage shared storage systems, that is, storage 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 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
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] |
```

#### Prerequisites

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

**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 on [page 319](#).

**-gname** *StorageGroupName*

Specifies the name of the storage group.

**-uid** *StorageGroupUID*

Specifies the storage group unique identification number.

For Fibre Channel storage 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 storage 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 storage-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 storage-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 storage system.

**-alu** *alunumber*

Specifies the storage-system LUN number, the actual LUN number internal to the storage 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*

**-addsnapshot** adds a SnapView 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 storage 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 storage 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 storage-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 storage 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 storage 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 on [page 319](#).

**-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 on [page 319](#).

**-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** *hlunnumber*

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** *hlunumber*

*hlunumber* specifies the logical unit number, which is the number that the storage-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* | **-snapshotwwn** *snapshotUID*

**-removesnapshot** removes a SnapView 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*]  
**-arraycomppath** **0** | **1**] [**-failovermode** **0** | **1** | **2** | **3** | **4**]  
**-type** *initiator-type*] [**-unitserialnumber** **array** | **lun**] [**-o**]

The **-sethost** function creates or changes certain settings for all initiators (HBAs) in the specified host that are connected to a specific storage 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 storage systems, the **storagegroup** command with the **-sethost** switch will accept nondefault values for the following switches: **-failovermode**, **-arraycomppath**, **-type**, and **-unitserialnumber**.

**-arraycommpath [ 0 | 1 ]**

**storagegroup -sethost** with **-arraycommpath** modifies one or more communication paths between the server and storage system. It may be needed when there is no LUN owned by an SP through which the host recognizes the storage system and in the storage group through which the host recognizes the storage system. A communication path is a logical path between the host and the storage system.

---

Note: AX series storage systems support only **arraycommpath** setting 1.

---

Changing the **arraycommpath** setting may force the storage system to reboot. The CLI will issue a confirmation request that indicates whether a storage-system reboot is required. However, the CLI will not issue a confirmation request first 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 storage 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 ss1_spa storagegroup -sethost -hostname
server101 -arraycommpath 1
```

**storagegroup -sethost** with **-failovermode** sets the type of trespass needed for failover software. Use this command with certain types of failover software if suggested in the failover software Release Notice or manual.

Changing the **failovermode** setting may force the storage system to reboot. The CLI will issue a confirmation request that indicates whether a storage-system reboot is required. However, the CLI will not issue a confirmation request first 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 storage systems with storage group only. For storage systems without storage group, use the command **failovermode** (see [Chapter 3, "Basic Commands"](#)).

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 the EMC Powerlink website.

**0**  
Failover mode 0.

**1**  
Failover mode 1.

**Note:** AX series storage systems support only failover mode setting 1.

**2**  
Failover mode 2.

**3**  
Failover mode 3.

## 4

Failover mode 4.

Example with **failovermode**:

```
naviseccli -h ss1_spa storagegroup -sethst -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 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**type*

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  CLARiiON Open (default).  
0x9  SGI  
0xa  HP (no auto trespass)  
0x13 Dell  
0x16 Fujitsu  
0x1c Compaq Tru64
```

**Note:** AX series storage systems support only type 0x3.

**-unitserialnumber** [*array* | *lun*]

The **unitserialnumber** command displays or sets the VDPP80 behavior mode, which defines how subsequent VPDP80 requests will be populated with the storage-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 -sethst** above.

**array**

Sets the behavior mode to storage system.

**lun**

Sets the behavior mode to LUN.

---

**Note:** Not supported on AX series storage systems.

---

Example with **unitserialnumber**:

```
naviseccli -h ss1_spa storagegroup -sethost
-unitserialnumber lun
```

**storagegroup** command, **-setpath** switch

```
-setpath [-gname StorageGroupName | -uid StorageGroupUUID]
[-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.

---

**Note:** The **storagegroup** command with the **-setpath** switch is supported on AX4-5 series storage systems, but is not supported on other AX series storage systems.

---

**-arraycommpath 0 | 1**

**storagegroup -setpath** with **-arraycommpath** modifies one or more communication paths between the server and storage system. It may be needed when there is no LUN owned by an SP through which the host recognizes the storage system and in the storage group through which the host recognizes the storage system.

Changing the **arraypath** setting forces the storage system to reboot.

This form of **-arraycommpath** works for storage systems with storage group only. For storage systems without storage group, use the command **arraycommpath** (see “[arraycommpath](#)” on page 59).

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 if suggested in the *Installation Roadmap for CLARiiON Storage Systems*, which is on the EMC Powerlink website.

Changing the **failovermode** setting may force the storage system to reboot. The CLI will issue a confirmation request that indicates whether a storage-system reboot is required. However, the CLI will not issue a confirmation request first if you specified the **-o** switch.

This form of **-failovermode** works for one storage group. To set the entire storage-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 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 -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 storage system uses the hostname (or **-ip** *address* switch) for any subsequent **storagegroup -connecthost** and **-disconnecthost** commands you issue. If you omit this, the storage system uses the HBA UID. Not required.

**-ip** *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 unitserialnumber 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** indicate whether the storage group can be shared. Only Navisphere Manager 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

In the following example, Navisphere connects all of **ss1\_spa**'s HBA-SP combinations to the storage group **mail**.

```
naviseccli -h ss1_spa storagegroup -gname mail -connecthost -o
```

In the next, 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

The following example 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
```

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 -hbuid xxxxx-sp a -sport xxxxx -failovermode 4
```

The following example specifies the form of failovermode which works 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 storage-system type

#### Description

The **naviseccli systemtype** command changes the configured storage-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 storage system running storage group. Do not change the system type of a storage system unless it has been quiesced. The agent must be running and managing the storage system.

Depending on the type of HP-UX operating system, you may want auto-trespass to be on or off.

#### Syntax

**systemtype** is used with **naviseccli** (described on [page 41](#)) as follows:

```
systemtype -config [type] [-o]
```

#### Prerequisites

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

#### Options

**-config [type]**

Sets the storage-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** - CLARiiON open storage system (default)

**9** - SGI (CX600)

**13** - Dell

**16** - Fujitsu

**1C** - HP type

**-o**

Executes the command without prompting for confirmation.

**Examples** `naviseccli -h ss1_spa systemtype`

The configured system type is: 0x3

The current system type is: 0x3

**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:

**naviseccli -h ss1\_spa setcache -wc 0.**

This operation will cause a storage system reboot!

DO YOU WISH TO CONTINUE? (y/n) **y**

This example sets the storage-system type to an HP-UX system type with auto-trespass on. The newly configured system type will become the current system type after the SP restarts.

## trespass

### Trespasses one or more LUNs to control of an SP

**Description** This command trespasses one or more LUNs to control of an SP.

**Note:** Not applicable to single-SP storage 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** is used with **naviseccli** (described on [page 41](#)) as follows:  
**trespass** [**all** | **lun** *lun-number* | **mine**]

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

**Options**

- all**  
Trespasses on all LUNs in the system.
- lun** *lun-number*  
Trespasses on one LUN (specified in *lun-number*).
- mine**  
Performs whatever trespassing is necessary to take all of the LUNs that have this SP as a default owner.

**Output** There is no output. Errors are printed to **stderr**.

**Example** `naviseccli -h ss1_spa trespass all`

## unbind

### Deconfigures a LUN

**Description** The **naviseccli unbind** command deconfigures LUNs from their current logical configuration. **unbind** 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
unbind lun [-o]
```

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

**Options** *lun*

Specifies the LUN to deconfigure.

**-o**

Executes the command without prompting for confirmation.

**Examples** **naviseccli -h ss1\_spa unbind 3 -o**

This command destroys (deconfigures) LUN 3 without prompting you for permission, and frees its disks to be reconfigured.

```
naviseccli -h ss1_spa unbind 0 1 2 3 4 -o
```

This command destroys (deconfigures) LUNs 0, 1, 2, 3, and 4 without prompting you for permission, and frees its disks to be reconfigured.

## unitserialnumber

**Displays or changes the VDPP80 (port 80) behavior mode**

**Description** The **naviseccli unitserialnumber** command displays or set the VDPP80 behavior mode, which defines how subsequent VPDP80 requests will be populated with the storage-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 [page 305](#).

**Syntax** **unitserialnumber** is used with **naviseccli** (described on [page 41](#)) as follows:

```
unitserialnumber [array | lun] [-o]
```

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

**Options** **array**

Sets the behavior mode to array (storage system).

**lun**

Sets the behavior mode to LUN.

**-o**

Executes the command without prompting for confirmation.

**Examples** **naviseccli -h ss1\_spa unitserialnumber**

```
Current unit serial number mode is: LUN
```

```
naviseccli -h ss1_spa unitserialnumber array
```

These two commands display the current behavior mode and then set the mode to **array** (storage-system) mode.















## LUN Expansion (MetaLUN) Commands

This chapter explains the Navisphere Manager CLI commands that manage Navisphere metaLUNs.

Major sections in the chapter are:

- ◆ MetaLUNs overview ..... 328
- ◆ CLI metaLUN command functions ..... 330
- ◆ metalun -destroy ..... 331
- ◆ metalun -expand ..... 332
- ◆ metalun -info ..... 341
- ◆ metalun -list ..... 343
- ◆ metalun -modify ..... 350
- ◆ Using the added capacity of an expanded SnapView, MirrorView or SAN Copy LUN ..... 352

## 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 (base LUN) into a larger unit called a metaLUN. You do this by adding LUNs to the base 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 storage system.

A metaLUN may 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 which includes the base 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.



Figure 3 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 base LUN re-striped across all three LUNs.




Figure 3 Sample stripe expansion

Figure 4 depicts a LUN using concatenate expansion to add two LUNs. The result is a metaLUN with two components. The first component is the base LUN. The second component consists of two new LUNs striped together, but not striped to the first component.




Figure 4 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 the following table.

**Table 1** **naviseccli metaLUN command functions, listed in order of use**

<b>Essential functions (in order performed)</b>	
<b>metalun -expand</b>	Creates a metaLUN or expands one
<b>metalun -destroy</b>	Destroys a metaLUN and then unbinds all LUNs that were part of the metaLUN
<b>Optional and status functions (alphabetically)</b>	
<b>metalun -info</b>	Provides metaLUN information with respect to a specific storage system
<b>metalun -list</b>	Lists the properties of existing metaLUNs
<b>Optional reconfiguration functions (alphabetically)</b>	
<b>metalun -modify</b>	Modifies certain metaLUN properties

## metalun -destroy

**Destroys a metaLUN and unbinds all LUNs that were part of the metaLUN**

**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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
metalun -destroy -metalun number | WWN | [-o]
```

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

**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 343](#)).

**-o**

Executes the command without prompting for confirmation.

**Example** **naviseccli metalun -destroy 20**

This command destroys the metaLUN with the number 20.

## metalun -expand

### Description

Creates a new metaLUN or expands the capacity of an existing one

The `naviseccli metalun` command with the `-expand` function creates a new metaLUN or expands the capacity of an existing one. See [“Requirements and recommendations” on page 332](#) for metaLUN requirements.

---

**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 Powerlink.

---

### Requirements and recommendations

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 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.

#### 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.

---

### Syntax

**metalun -expand** is used with **naviseccli** (described on [page 41](#)) to create a *new* metaLUN or expand an *existing* metaLUN. There are two distinct versions of this command, one for Navisphere Manager 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 Navisphere Manager and requires you to bind the destination LUN before running the command.

---

**Note:** The following command syntax works only with Navisphere Manager.

---

To create a new metaLUN, the following command switches are valid:

```
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 ]
```

To expand a metaLUN, the following command switches are valid:

```
metalun - expand -base number | WWN -lus number | WWN [-type
C | S ] [-sq tb | gb | mb | bc ] [-cap number ] [-nogrow ] [-o
```

**Note:** The following command syntax works only with Navisphere Express.

To create a new metaLUN and expand an existing metaLUN for a storage system running Navisphere Express, you must use this command.

```
metalun -expand -base number | WWN  
[-sq tb | gb | mb -cap number] | [-percent number] [-o]
```

#### Prerequisites

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

#### Options

The command switches for Navisphere Manager (See [page 334](#)) and Navisphere Express (See [page 339](#)) are given below:

The following command switches are supported with Navisphere Manager:

**-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 343](#)).

When you specify a LUN with the **-base** command, the CLI creates a new metaLUN with the number and WWN of the base LUN. Once the CLI creates the new metaLUN, it assigns a new WWN to the original base LUN and makes it a private LUN. When expanding a single LUN, you can use any or all of the **-expand** switches.

**Note:** Ú The base 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 343).

---

**Note:** Expansion LUNs cannot be participating in a SnapView (snapshot or clones), MirrorView, or SAN Copy session.

---

**-autoassign *0* | *1*** (Not supported on AX4-5 series systems running Navisphere Manager)

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 base LUN.

Autoassign controls the ownership of the LUN when an SP fails in a storage 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 storage 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 storage 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*** (Not supported on AX4-5 series systems running Navisphere Manager)

---

**Note:** Available only if the **preservedata** switch is set to **0**.

---

EMC does not recommend using a metaLUN **bindoffset** value if the base LUN already has one.

---

Automatically aligns the logical sector number of the metaLUN to a stripe boundary on the base 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 `bindoffset` 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 352.

---

**-defaultowner** A | B

Specifies the SP that will assume ownership of the metaLUN when the storage system is powered up. If you do not specify a default owner, it defaults to the owner of the base LUN.

**-elszm** *elementSizeMultiplier* (Not supported on AX4-5 series systems running Navisphere Manager)

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. For more information on the metaLUN stripe multiplier, refer to the *EMC CLARiiON Best Practices for Fibre Channel Storage: FLARE Release 26 Firmware Update*, available on the Powerlink website.



**-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.

**-name** *metaLUN-name*

Specifies the name of the metaLUN. If you do not specify a name, it defaults to the name of the base 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.

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 [page 338](#)). Allows you to preserve or overwrite any existing data on the base LUN, only if the base LUN is outside a storage group.

---

**Note:** You cannot use this command with the **-type C** switch.

---

If the base LUN is outside of a storage group, do one of the following:

- Specify **1** to enable stripe expansion and preserve the data on the base LUN. Note that this operation could take a while. The default value is **1**.

---

**Note:** You must *always* specify **1** when the base LUN is in a storage group. 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 base 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. The default qualifier is block count.

---

**Note:** ÚXXXXXXXXXXXXXXXXXXXXX AX4-5 series storage systems running Navisphere Manager support only terabytes (TB), gigabyte (GB), and megabyte (MB) size qualifiers.

AX series systems support only megabyte (MB) and gigabyte (GB) size qualifiers.

---

**-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 base 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 [“Requirements and recommendations” on page 332](#).

If you are expanding a LUN (the base LUN), the **-type S** switch creates a metaLUN with one component which includes the base LUN and any LUNs that you add. Any data on the base LUN is restriped across the base LUN and the new LUNs.

---

**Note:** If the base LUN has existing data, and it is outside of a storage group, refer to the **-preservedata** switch ([page 337](#)) for additional options.

---

You can expand only the last component of an existing metaLUN. In this case, the **-type S** switch restripes 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.

The following command switches are supported with Navisphere Express and are mandatory:

**-base** *number* | *WWN*

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*

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*

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.

The optional switch for Navisphere Express is:

**-percent** *number*

Specifies the amount of additional storage you want to expand by, as a percent 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.

---

**Note:** For AX4-5 series systems running Navisphere Express, the default expansion type is concatenate. The expansion type, stripe is not supported on AX4-5 series systems.

---

The AX4-5 series systems support only an expansion rate of **high**.

**Example Navisphere Manager:**

```
naviseccli metalun -expand -base 20 -lus 42 -name newMetaLUN
```

This command creates a metaLUN with the name **newMetaLUN**, which includes base LUN 20 and LUN 42. Since the default expansion type is stripe, this metaLUN is striped.

**Navisphere Express:**

```
naviseccli metalun -expand -base 1 -sq gb -cap 1
```

This command expands a metaLUN by 1 GB.

## metalun -info

**Provides general information about a storage system's metaLUN environment**

**Description** The `naviseccli metalun` command with the `-info` function and no other options displays all general information about the metaLUN environment for a storage system. Issuing this command with one or more options displays the information based on the options issued.

**Syntax** `metalun -info` is used with `naviseccli` (described on [page 41](#)) as follows.

`metalun -info [-canexpand] [-nummetaluns] [-maxmetaluns] [-maxcomponents] [-maxluns] [-availableluns] [-lunsinmeta]`

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

**Options** `-canexpand`

Specifies whether you can or cannot expand any LUNs or metaLUNs on the storage system. - **Yes** or **No**.

`-nummetaluns`

Displays the number of metaLUNs in the storage system.

`-maxmetaluns`

Displays the maximum number of metaLUNs that the storage system will support (based on storage-system type). See the CLI release notes for currently supported values.

`-maxcomponents`

Displays the maximum number of components allowed per metaLUN for this storage system (based on storage-system type). See the CLI release notes for currently supported values.

`-maxluns`

Displays the maximum number of LUNs per component that this storage system supports (based on storage-system type). See the CLI release notes for currently supported values.

**-availableluns**

Lists all LUNs (by LUN number) in the storage system that are available for expansion.

**-lunsinmeta**

Lists any LUNs (by LUN number) in the storage system that are part of a metaLUN.

**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.
```

**Example**

**naviseccli -address 10.14.20.57 metalun -info**

This command requests information about the metaLUN environment for the storage system with SP 10.14.20.57.

## metalun -list

### Lists the properties of existing metaLUNs

**Description** The **naviseccli metalun** command with the **-list** function and no other options lists all property information for all metaLUNs on a storage 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 storage system. Use Navisphere CLI or Navisphere Manager. 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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
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] [-totalcap] [-whist*]
```

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

**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 storage 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 storage 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 storage system.

**-aac\***

Displays the active/active counters associated with the metaLUN.

**-actualusercap**

Displays the current user capacity of all metaLUNs on the storage 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 storage 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 storage system to a stripe boundary on the base LUN (see [page 335](#)).

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 storage 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 storage 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 storage 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**

Displays the ability of all metaLUNs on the storage system to be expanded - **Yes** or **No**.

Issuing the **-canexpand** command with a metaLUN name, number or WWN displays the ability of the specified metaLUN to be expanded - **Yes** or **No**.

**-defaultowner**

Lists the default SP owner of all metaLUNs on the storage 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 storage 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 storage system (see [page 337](#)). 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 [page 337](#)).

**-elszm**

Displays the element size multiplier for all metaLUNs on the storage system (see [page 336](#)).

Issuing the **-elszm** command with a metaLUN name, number or WWN displays the element size multiplier for the specified metaLUN (see [page 336](#)).

**-isredundant**

Displays the redundancy of all metaLUNs on the storage system - **Yes** if it is redundant; **No** if it is not redundant.

Issuing the **-elszm** command with a metaLUN name, number or WWN displays the redundancy of the specified metaLUN - **Yes** if it is redundant; **No** if it is not redundant.

**-nzcra**

Returns the nonzero request count arrivals in SP A and SP B.

**-percentexp**

Displays the percent complete of the stripe expansion process of all base LUNs or metaLUNs on the storage system.

Issuing the **-percentexp** command with a metaLUN name, number or WWN displays the percent complete of the stripe expansion process of the specified base LUN or metaLUN.

**-rhist\***

Returns read histogram information for all metaLUNs on the storage system.

---

**Note:** A read histogram is an array of 10 locations that contain the number of reads. Element  $n$  of the array 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 storage 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 storage system.

When issuing the **-state** command with a metaLUN name, number or WWN displays the state of the specified metaLUN.

Table 2 Valid metaLUN states

State	Description
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 (restripping data across all component LUNs).

**-totalcap**

Displays the total available capacity of all metaLUNs on the storage system.

Issuing the **-totalcap** command with a metaLUN name, number or WWN displays the total available capacity of the specified metaLUN.

**-whist\***

Returns write histogram information for all metaLUNs on the storage system.

**Note:** A write histogram is an array of 10 locations that contain the number of writes. Element  $n$  of the array 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.

### 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

```

**Example** **naviseccli -address 10.14.20.57 metalun -list**

This command requests information on all metaLUNs that are currently owned by the SP with IP address 10.14.20.57.

## metalun -modify

**Modifies certain properties of an existing metaLUN**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
metalun -modify -metalun number | WWN [-autoassign 0 | 1]
[-capcapacity] [-defaultowner A | B]
[-expansionrate low | medium | high | ASAP | value]
[-newname newName] [-o] [-sq tb | gb | mb | bc] [-tomaxcap]
```

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

**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 base LUN. For more information about auto assign, see [page 335](#).

`-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 352](#).

---

`-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 storage 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.

---

**Note:** AX4-5 series storage systems running Navisphere Manager support only terabyte (TB), gigabyte (GB), and megabyte (MB) size qualifiers.

---

**-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 “[naviseccli metalun -destroy 20](#)” on page 331), and to allow the host to access the added capacity.

---

**Note:** For a stripe expansion, any existing data is restriped across the original LUN or metaLUN and the added LUNs *during* the SnapView, MirrorView or SAN Copy session. This restriping is invisible to the host.

---

### For LUNs in a SnapView session

To access the added capacity of an expanded snapshot source LUN, do the following:

1. Stop *all* SnapView sessions and destroy *all* snapshots associated with the expanded source LUN (see the *admsnap*, the online help).
2. Use the **metaLUN -modify** command on [page 350](#) to change the capacity of the expanded LUNs.
3. Make sure the host operating system can utilize the added capacity.
4. 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:

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** command on [page 350](#) 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:

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** command on [page 350](#) 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:

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** command on [page 350](#) 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 chapter 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.

Major sections are:

- ◆ Introduction to the reserved LUN pool ..... 356
- ◆ Estimating the reserved LUN pool size ..... 357
- ◆ reserved -lunpool -addlun ..... 359
- ◆ reserved -lunpool -list ..... 361
- ◆ reserved -lunpool -rmlun ..... 364

---

## Introduction to the reserved LUN pool

The global reserved LUN pool works with replication software, such as SnapView, SAN Copy, and MirrorView/A to store data or information required to complete a replication task. The reserved LUN pool consists of one or more private LUNs. The LUN becomes private when you add it to the reserved LUN pool. Since the LUNs in the reserved LUN pool are private LUNs, they cannot belong to storage groups and a server cannot perform I/O to them.

Before starting a replication task, the reserved LUN pool must contain at least one LUN for each source LUN that will participate in the task. You can add any available LUNs to the reserved LUN pool. Each storage system manages its own LUN pool space and assigns a separate reserved LUN (or multiple LUNs) to each source LUN.

All replication software that use the reserved LUN pool shares the resources of the reserved LUN pool. For example, if you are running an incremental SAN Copy session on a 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.

---

## Allocation of the reserved LUN pool

The software allocates reserved LUNs on a per source LUN basis. Each storage system manages its own reserved LUN pool space and allocates reserved LUNs on a per-source LUN basis, not a per-session basis. For example, if two replication sessions are running on a source LUN, the storage system assigns reserved LUNs, as needed, to the source LUN, and both sessions share these LUNs.

The assignment of reserved LUNs is based upon the first available free LUN in the global reserved LUN pool. An adequate number of reserved LUNs is essential because the software will terminate sessions if no free LUNs are available in the global reserved LUN pool. If you are running multiple sessions per source LUN and your reserved LUN fills up (runs out of space) and no free LUNs are available in the global reserved LUN pool, the session that is trying to allocate a reserved LUN from this pool will terminate.

---

## Estimating the reserved LUN pool size

Each reserved LUN can vary in size. However, using the same size for each LUN in the pool is easier to manage because the LUNs are assigned without regard to size; that is, the first available free LUN in the global reserved LUN pool is assigned. Since you cannot control which reserved LUNs are being used for a particular replication session, we recommend that you use a standard size for all reserved LUNs.

If you want to optimize space utilization, the recommendation would be to create many small reserved LUNs, which allows for sessions requiring minimal reserved LUN space to use one or a few reserved LUNs, and sessions requiring more reserved LUN space to use multiple reserved LUNs. On the other hand, if you want to optimize the total number of source LUNs, the recommendation would be to create many large reserved LUNs, so that even those sessions which require more reserved LUN space only consume a single reserved LUN.

---

## Estimating a suitable reserved LUN pool size

The following considerations should assist in estimating a suitable reserved LUN pool size for the storage system.

- ◆ If you wish to optimize space utilization, use the size of the smallest source LUN as the basis of your calculations (see [page 358](#)). If you wish to optimize the total number of source LUNs, use the size of the largest source LUN as the basis of your calculations (see [page 358](#)).
- ◆ If you have a standard online transaction processing configuration (OLTP), use reserved LUNs sized at 10-20%. This tends to be an appropriate size to accommodate the copy-on-first-write activity.
- ◆ If you plan on creating multiple sessions per source LUN, anticipate a large number of writes to the source LUN, or anticipate a long duration time for the session, you may also need to allocate additional reserved LUNs. With any of these cases, you should increase the calculation accordingly. For instance, if you plan to have 4 concurrent sessions running for a given source LUN, you might want to increase the estimated size by 4 – raising the typical size to 40-80%.

---

**Note:** To protect your replication sessions from being terminated (if no free reserved LUNs were available in the pool), the number of reserved LUNs in the pool should always be two times the number of source LUNs.

---

**Example for optimizing for capacity**

The goal is to create 1 session each on 100 source LUNs, where the largest is 100 GB, and the smallest is 50 GB:

**Basis of source LUN size:** 50 GB

**Estimated change rate:** 20% = 10 GB

**Minimum number of reserved LUNs:** 100

**Extra reserved LUNs to ensure that larger source LUNs get enough reserved LUN space:** 2 x minimum

This would translate to 200 reserved LUNs that are each 10 GB

**Example for optimizing for max source LUNs**

The goal is to create 4 sessions each on 100 source LUNs, where the largest is 100 GB, and the smallest is 50 GB:

**Basis of source LUN size:** 1 TB

**Estimated change rate:** 4 x 20% = 800 GB

**Minimum number of reserved LUNs =** 100

**Extra reserved LUNs to ensure that larger source LUNs get enough reserved LUN space:** + 20% = 20 extra reserved LUNs

This would translate to 120 Reserved LUNs that are each 80 GB

## reserved -lunpool -addlun

### Adds a LUN to the reserved LUN pool

**Description** The **naviseccli reserved** command with the **-lunpool** and **-addlun** functions adds one or more LUNs to the reserved LUN pool. Each storage 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 storage 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** is used with **naviseccli** (described on [page 41](#)) as follows:

**reserved -lunpool -addlun** *lun\_numbers*

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

**Options** *lun\_numbers*

Specifies the logical unit number(s) to add to the reserved LUN pool.

**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**.

**Example** `naviseccli -h ss1_spa reserved -lunpool -addlun 11`

For `ss1_spa`, this command adds the LUN with ID 11 to the storage system's reserved LUN pool.



## reserved -lunpool -list

### Displays reserved LUN pool information

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`reserved -lunpool -list [-allocatedluns] [-chunksize] [-freeluncount] [-freeluns] [-freesize] [-luncount] [-percentused] [-size] [-used]`

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

**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 storage 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.

**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 storage system is running Navisphere CLI version 6.24 or later, reserved LUN pool information will display for the entire storage system, regardless of the SP you direct the command to. If your storage system is running Navisphere CLI version 6.22 or earlier, reserved LUN pool information will display 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
Target LUN:                                   170
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`.

**Example** `naviseccli -h ss1_spa reserved -lunpool -list`

This command displays the storage system's reserved LUN pool information. A sample output follows.

## reserved -lunpool -rmlun

**Removes one or more LUNs from the reserved LUN pool**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
reserved -lunpool -rmlun luns [-o]
```

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

**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.

**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`.

**Example** `naviseccli -h ss1_spa reserved -lunpool -rmlun 42`

For `ss1_spa`, this command starts removing LUN 42 from the storage system's reserved LUN pool. The software asks for confirmation:

```
Do you really want to remove the luns from LUN pool.  
(y/n) [N] ?
```

```
y
```

This chapter explains each of the **naviseccli** iSCSI(Internet SCSI) commands; that is, the CLI commands that are common to all CLARiiON iSCSI disk-array storage systems.

Major sections are:

◆ About iSCSI storage systems in IP networks.....	367
◆ Using iSNS within an iSCSI storage environment .....	370
◆ iSCSI commands overview .....	372
◆ connection -adduser .....	374
◆ connection -deleteuser .....	376
◆ connection -getuser .....	378
◆ connection -getport .....	380
◆ connection -setport .....	384
◆ connection -delport .....	386
◆ connection -route.....	387
◆ connection -pingnode.....	389
◆ connection -traceroute .....	391
◆ connection -setshredauth .....	393
◆ connection -delshredauth .....	395
◆ connection -getshredauth .....	396
◆ connection -addset .....	397
◆ connection -delset .....	399
◆ connection -modifyset .....	400
◆ connection -getset.....	402
◆ connection -addpath.....	404
◆ connection -delpath .....	406
◆ connection -modifypath.....	407
◆ connection -verifypath .....	409
◆ isns -addserver .....	411


- ◆ isns -deleteserver ..... 413
- ◆ isns -listserver..... 415
- ◆ isns -setprimary..... 416

## About iSCSI storage systems in IP networks

An Internet SCSI (iSCSI) storage system, such as the CX500i, connects directly to an IP network. Servers with either iSCSI HBAs (host bus adapters) or Ethernet connections can access the iSCSI storage system through an IP network, such as a local area network (LAN).

A LAN is a set of point-to-point connections between nodes, with each node having its own unique IP address. Connections are made through one or more network components such as switches or hubs. Nodes are connected through a LAN by Ethernet CAT 6 (for Gigabit Ethernet LAN) and CAT 5 (for 10/100 LAN) copper cables. Network switches are *not* nodes.

Each node in an iSCSI environment is either an initiator (server) or a target (storage system). [Figure 5](#) shows an initiator node and a target node.



**Figure 5** Nodes — Initiator and target

Before a server is able to initiate server I/O to the iSCSI storage system, it must be configured as follows:

- ◆ You have installed one of the following interface cards and relevant drivers:
  - Supported iSCSI HBA cards that have a driver and configuration tool (for example, QLogic), or
  - Gigabit Ethernet Network Interface Cards (NICs) running Microsoft software that provides HBA functionality.

---

**Note:** We support 10 Mb, 100 Mb, and 1000 Mb Gigabit Ethernet (GigE) interfaces, but the storage system supports only 1000 Mb. If your NIC does not run GigE, then you need to connect to the storage system using a GigE router or switch.

---

- ◆ You have cabled the storage system properly (refer to the setup guide that shipped with the storage system).
- ◆ You have installed the Navisphere Utilities on each server with access to data on the storage system (refer to the setup guide).
- ◆ You have installed PowerPath software on the servers for multi-pathing (refer to the setup guide).
- ◆ You have set the network parameters and security for the SP management ports on the storage system.

The iSCSI interface uses CHAP (Challenge Handshake Authentication Protocol) to protect the storage system's iSCSI ports from unwanted access. CHAP is *optional*, but if your storage system might be accessed from a public IP network, we strongly recommend that you use CHAP security.

CHAP is a method for authenticating iSCSI users (initiators and targets). The iSCSI storage system can use CHAP to authenticate server initiators and initiators can authenticate targets such as the storage system. To use CHAP security, you must configure CHAP credentials for the storage system iSCSI ports and any servers that will access the storage-system data.

---

**Note:** If you will be using CHAP security, we strongly recommend that you configure it on both the storage system and the server *before* initiating server I/O.

---



Figure 6 illustrates an iSCSI storage system in a storage domain with various servers and illustrates how iSCSI storage systems differ from Fibre Channel storage systems.

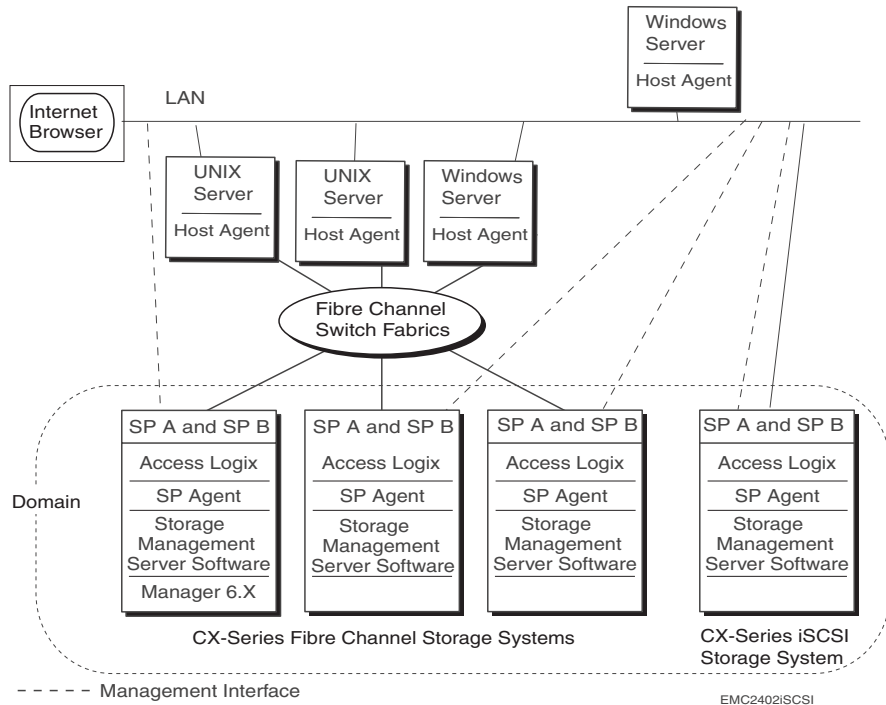


Figure 6 Fibre Channel and iSCSI storage systems in an IP network

## Using iSNS within an iSCSI storage environment

**Note:** iSNS (Internet Storage Naming Service) is supported only on Windows platforms that are part of an iSCSI network configuration.

The iSNS service provides the same function for TCP/IP storage networks as the Simple Name Server (SNS) service in a Fibre Channel fabric — automated discovery, management and configuration of iSCSI devices. It eliminates the need to manually configure each individual storage device with its own list of initiators and targets. Once configured, the iSNS server assumes responsibility for the discovery and management of iSCSI devices.

The iSNS service includes an iSNS server component and iSNS client components. The iSNS server must reside somewhere within the IP storage network, for example, in the switch firmware, or on a host. An iSNS client resides on both the iSCSI storage system and any iSCSI hosts connected to the storage system. When you start the storage system, the iSNS client on the storage system gathers all the storage system's iSCSI port information and stores it locally on the storage system. When you add a server to the storage system's iSNS configuration, Navisphere establishes a connection from the storage system to the iSNS server, and then registers all the stored information on the storage system with the iSNS server.

To use the features of the iSNS service in an iSCSI storage environment:

- ◆ All the iSCSI requirements listed in [“About iSCSI storage systems in IP networks” on page 367](#), must be met.
- ◆ An iSNS server must be running on a network to which the storage system has access.
- ◆ The host initiators or HBAs must be able to use the iSNS service;
- ◆ The storage system must support the iSNS service.
- ◆ The iSNS server must be able to communicate with the storage system.

Figure 7 represents a sample iSNS configuration. In this configuration, the management and storage networks are separate. The management station is used to configure both the iSNS servers and the storage system (using Navisphere Manager). The iSNS servers are also connected to the storage network so that the clients may query them for target information. The storage system uses either the storage or the management network to register its targets (ports) with the iSNS servers.




Figure 7 Sample iSNS storage configuration

## iSCSI commands overview

The **naviseccli** iSCSI commands in this chapter include the connection and iSNS commands. [Table 3](#) identifies the iSCSI commands that can be used.

**Table 3** **naviseccli iSCSI commands**

Command	Description
<b>CHAP User Commands</b>	
<b>connection -adduser</b>	Creates a CHAP user account.
<b>connection -deleteuser</b>	Deletes a CHAP user account.
<b>connection -getuser</b>	Displays CHAP user accounts.
<b>Port administration commands</b>	
<b>connection -getport</b>	Retrieves a specific iSCSI port or the list of all the iSCSI ports.
<b>connection -setport</b>	Modifies the properties of a connection port.
<b>connection -delport</b>	Deletes a virtual port.
<b>connection -route</b>	Sets up the static routes.
<b>connection -pingnode</b>	Verifies that packets reach a destination node without errors.
<b>connection -traceroute</b>	Displays the route that a packet takes to a destination node.
<b>Connection set database commands</b>	
<b>connection -setshredauth</b>	Sets a username and secret for the storage system's shared credentials record.
<b>connection -delshredauth</b>	Removes the shared authentication information from the system.
<b>connection -getshredauth</b>	Displays the username and the switch enabled for the shared authentication record.
<b>connection -addset</b>	Adds a connection set.
<b>connection -delset</b>	Deletes a connection set and associated connection paths.

Command	Description
<b>connection -modifyset</b>	Modifies parameters of the connection set.
<b>connection -getset</b>	Lists a specific connection set or all the connection sets.
<b>connection -addpath</b>	Adds a connection path to a specific connection set.
<b>connection -delpath</b>	Deletes a connection path from a connection set.
<b>connection -modifypath</b>	Modifies a connection path from a connection set.
<b>connection -verifypath</b>	Tests a connection path from a connection set.
<b>iSNS commands</b>	
<b>isns -addserver</b>	Adds an iSNS server to a storage system's iSNS service configuration.
<b>isns -deleteserver</b>	Deletes an iSNS server from a storage system's iSNS service configuration.
<b>isns -listserver</b>	Displays all iSNS servers that are configured in the storage system's iSNS service.
<b>isns -setprimary</b>	Sets a configured iSNS server as the primary server for a storage system's iSNS service configuration.

## connection -adduser

### Adds a CHAP authentication user

**Description** The **naviseccli connection** command with the **-adduser** function adds either an initiator CHAP authentication user or a target (storage system) CHAP authentication user to the storage-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** is used with **naviseccli** (described on [page 41](#)) as follows:

**connection -adduser -definedFor initiator | target -initiatorName initiatorName [-userName userName] -secret secret [-hexsecret] [-o]**

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

**Options** **-definedFor initiator | target**

**initiator** adds an initiator CHAP user account to the storage system user database. If you specify **initiator**, you must specify an initiator using the **-initiatorName** parameter.

**target** specifies that a target (storage 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 storage 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 storage 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

**Description** The `naviseccli connection` command with the `-deleteuser` function deletes a specified initiator or target CHAP authentication user from the storage system CHAP user database.

**Syntax** `connection -deleteuser` is used with `naviseccli` (described on [page 41](#)) as follows:

`connection -deleteuser -definedFor initiator | target -initiatorName initiatorName -userName username [-o]`

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

**Options** `-definedFor initiator | target`

**initiator** specifies that an initiator user account will be deleted from the storage-system user database. If you specify **initiator**, you must specify an initiator using the `-initiatorName` parameter.

**target** specifies that the target (storage 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.

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



**Example** The following example deletes a user account from the storage 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
```

## connection -getuser

### Display CHAP authentication user data

**Description** The `naviseccli connection` command with the `-getuser` function displays CHAP authentication user account data that is configured on the storage system. If no arguments are specified, all of the CHAP users defined on the storage 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -getuser [-definedFor initiator | target ] [-userName
userName] [-initiatorName initiatorName]
```

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

**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 storage 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 storage 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**.

**Output**

```
Initiator Name:  iqn.1991-05.com.microsoft:cpc7745
User Name:      Guest
Defined For:    initiator
```

**Example** The following example gets information on the current user:  
`naviseccli connection -getuser`

## connection -getport

Retrieve a specific iSCSI port or the list of all the iSCSI ports

**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 storage system are displayed.

**Syntax** `connection -getport` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -getport [-sp a | b] [-portid portid [-vportid vportid] ]
[-vlanid] [-address IPAddress] [-subnetmask] [-gateway]
[-initiatorAuthentication] [-mtu] [-all]
```

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.

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

**Options** `-sp a | b`

Specifies SP A or B. Displays the information only for iSCSI ports configured on the specified SP.

`-portid portid`

Specifies an iSCSI 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 *dotted-quad* format). For example, 165.152.36.35.

**-subnetmask**

Displays the network mask for the specified iSCSI ports. The subnet mask uses the *dotted-quad* format. For example, 255.255.255.0.

**-gateway**

Displays the network gateway address that the specified iSCSI ports use. The gateway address uses the *dotted-quad* format. 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 storage system. **False** indicates that initiator authentication has been disabled for the port.

**-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.

**-all**

Displays all the physical and virtual port information for each iSCSI port.

**Output**

```
SP: A
Port ID: 1
Port WWN: iqn.1992-04.com.emc:cx.apm00034901792.a1
IP Address: 10.14.80.109
Subnet Mask: 255.255.255.0
Gateway Address: 10.14.80.1
Initiator Authentication: false
```

```
SP: B
Port ID: 0
Port WWN: iqn.1992-04.com.emc:cx.apm00034901792.b0
IP Address: 10.14.80.110
Subnet Mask: 255.255.255.0
Gateway Address: 10.14.80.1
```

```
Initiator Authentication: false
```

```
SP: B
Port ID: 1
Port WWN: iqn.1992-04.com.emc:cx.apm00034901792.b1
IP Address: 10.14.80.111
Subnet Mask: 255.255.255.0
Gateway Address: 10.14.80.1
Initiator Authentication: false
```

```
SP: A
Port ID: 0
Port WWN: iqn.1992-04.com.emc:cx.apm00034901792.a0
IP Address: 10.14.80.108
Subnet Mask: 255.255.255.0
Gateway Address: 10.14.80.1
Initiator Authentication: false
```

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: 0
Port WWN: iqn.1992-04.com.emc:cx.fnm00081500003.a0
iSCSI Alias: 0003.a0
Port Speed: 1000 Mb
Auto-Negotiate: Yes
Available Speeds: 10 Mb
-                : 100 Mb
-                : 1000 Mb
-                : Auto
Current MTU: 1500
Available MTU Sizes:
"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: false
```

```
Virtual Port ID: 1
VLAN ID: 202
IP Address: 192.168.51.140
Subnet Mask: 255.255.224.0
Gateway Address: 192.168.32.1
Initiator Authentication: false
```

- Examples**
- connection -getport** displays information about all iSCSI ports configured on the storage system, or nothing if no iSCSI ports are configured.
  - connection -getport -sp a -portid 0** displays information about port 0 on SP A, or prints an error message if the iSCSI port does not exist.
  - connection -getport -sp a** displays information about all iSCSI ports on SP A, or nothing if no iSCSI ports are configured on that SP.
  - connection -getport -portid 0** displays information about all iSCSI ports with ID 0, regardless of what SP they are configured on, or nothing if no matching ports are found.

## connection -setport

### Modify the properties of a port

<b>Description</b>	The <b>naviseccli connection</b> command with the <b>-setport</b> function modifies the properties of a specified SP port.
<b>Syntax</b>	<p><b>connection -setport</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a>) as follows:</p> <pre><b>connection -setport -sp a   b -portid portid [-vportid vportid] [-vlanid vlanid   -vlandisable] -address IPAddress -subnetmask subnetmask -gateway gateway [-initiatorAuthentication 0   1] [-mtu mtu] [-o]</b></pre>
<b>Prerequisites</b>	You must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<p><b>-sp a   b</b></p> <p>Specifies that the iSCSI port is on either SP A or SP B.</p> <p><b>-portid portid</b></p> <p>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.</p> <p><b>-vportid vportid</b></p> <p>Specifies the virtual port ID. If you do not specify the <b>-vportid</b>, the system defaults to virtual port 0.</p> <p><b>-vlanid vlanid</b></p> <p>Specifies the VLAN ID.</p> <p><b>-vlandisable</b></p> <p>Disables VLAN tagging on the virtual port.</p> <p><b>-address IPAddress</b></p> <p>Specifies the IP address of the iSCSI port. The address uses the <i>dotted-quad</i> format. For example: 165.152.36.35.</p> <p><b>-subnetmask subnetmask</b></p> <p>Specifies the subnetwork mask for the iSCSI port. The subnet mask uses the <i>dotted-quad</i> format. For example: 255.255.255.0.</p>



**-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 into the port must be authenticated by the storage system.

**0** = Not required - default

**1** = Required

**-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.

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

**Example** 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
```

```
Do you really want to perform this action (y/n)? y
```

## connection -delport

### Deletes a virtual port

<b>Description</b>	<p>The <b>naviseccli connection</b> command with the <b>-delport</b> function lets you delete a virtual port from a physical iSCSI port. To execute this command, more than one virtual port should be associated with the management port.</p> <p>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.</p>
<b>Syntax</b>	<p><b>connection -delport</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a>) as follows:</p> <pre><b>connection -delport -sp a   b -portid portid -vportid vportid [-o]</b></pre>
<b>Prerequisites</b>	For <b>naviseccli</b> , you must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<p><b>-sp a   b</b></p> <p>Deletes the virtual port of the specified SP. The default is the connected SP.</p> <p><b>-portid portid</b></p> <p>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.</p> <p><b>-vportid vportid</b></p> <p>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.</p> <p><b>-o</b></p> <p>Executes the command without prompting for confirmation.</p>
<b>Output</b>	Provides output depending on the options used.

## connection -route

Sets up the static routes to specific hosts or networks

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
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]
```

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

**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

**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 -tracert` commands to verify and diagnose network connectivity.

**Syntax** `connection -pingnode` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -pingnode [-sp a | b] [ [-portid portid -vportid vportid]
| -source sourceaddress] -address address [-packetSize packetSize]
[-count count] [-timeouts timeouts] [-delays delays]
```

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

**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.

**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
```

**Example**

```
naviseccli -h peregrine_spa connection -pingnode -sp a -portid 0
-vportid 0 -address 128.222.132.100
```

## connection -traceroute

**Display the route that a packet takes to a destination node**

**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 storage system iSCSI port to the target node. Use the **connection -traceroute** and **connection -pingnode** commands to verify and diagnose network connectivity.

**Syntax** **connection -traceroute** is used with **naviseccli** (described on [page 41](#)) as follows:

```
connection -traceroute [-sp a | b] -address IPAddress
[ [-portid portid -vportid vportid] | -source sourceaddress
[-mtu | -packetSize packetSize] ]
[-timeouts timeouts] [-delays delays]
```

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

**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.

**Output**

```
1: 172.20.1.1 1 ms 1500 mtu
2: 10.5.1.250 1 ms 1500 mtu
```

**Example**

The following example displays the route taken by a packet:

```
naviseccli -user a -password a -scope 0 -h ss1_spa connection  
-traceroute -sp a -address -mtu -source
```



## connection -setsharedauth

Sets username and secret for the storage system's shared credentials record

**Description** The `naviseccli connection` command with the `-setsharedauth` function sets a username and secret for the iSCSI initiator storage system's shared credentials record of the connection set database. The command lets you enable or disable the use of the storage system's shared authentication credentials.

**Syntax** `connection -setsharedauth` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -setsharedauth [-userName userName] [-secret secret
-hexsecret]] [-enable | -disable] [-o]
```

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

**Options** `-userName userName`

Specifies the CHAP username used by the initiator when challenged by the target storage 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 storage 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.

`-enable | -disable`

Enables or disables the use of the storage system's shared credentials.

**-o**

Executes the command without prompting for confirmation.

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

**Example** The following example sets a username and secret for the storage system's shared credentials record:

```
naviseccli connection -setsharedauth -username user1 -secret  
abcdefghi123 -enable -o
```

## connection -delsharedauth

**Removes the storage-system shared authentication credentials**

**Description** The `naviseccli connection` command with the `-delsharedauth` function removes the username and secret from the iSCSI initiator storage system.

**Syntax** `connection -delsharedauth` is used with `naviseccli` (described on [page 41](#)) as follows:

`connection -delsharedauth [-o]`

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

**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.

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

**Example** The following example removes the shared authentication credentials from the iSCSI initiator storage system:

`connection -delsharedauth -o`

## connection -getsharedauth

**Displays the username and status of the shared authentication for the iSCSI initiator storage system**

**Description** The `naviseccli connection` command with the `-getsharedauth` function displays the username and status of the shared authentication for the iSCSI initiator storage system.

**Syntax** `connection -getsharedauth` is used with `naviseccli` (described on [page 41](#)) as follows:

`connection -getsharedauth`

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

**Output** User Name: user1  
Enabled: True

If no shared credentials are defined for the storage system, then the output is:

No storage system shared credentials are defined.

**Example** The following example displays the username and status of the shared authentication for the iSCSI initiator storage system:

`connection -getsharedauth`

## connection -addset

### Adds a connection set

**Description** The `naviseccli connection` command with the `-addset` function adds a connection set to the connection set database.

**Syntax** `connection -addset` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -addset -name name [-digest none | Header | Data | Both] [-authpolicy none | shared | setspecific -userName userName -secret secret [-hexsecret]] [-o]
```

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

**Options** `-name name`

Specifies a unique identifier for the connection set. An arbitrary name supplied by the administrator, which 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 storage 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 storage 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 storage 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.

**Output**

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

**Example**

The following example adds a connection set to the connection set database:

```
connection -addset -name MirrorSet -authpolicy none -digest none -o
```

## connection -delset

**Deletes a connection set and associated connection paths**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`connection -delset -name name [-o]`

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

**Options** `-name name`

Specifies a unique identifier for the connection set. An arbitrary name supplied by the administrator, which 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.

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

**Example** The following example deletes a connection set and associated connection paths using the name of the connection:

`connection -delset -name MirrorSet -o`

## connection -modifyset

**Modifies parameters of a connection set**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
connection -modifyset -name name [-digest None | Header | Data | Both] [-authpolicy none | shared | setspecific -userName userName -secret secret [-hexsecret] ] [-o]
```

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

**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 storage 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 storage 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 storage 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.

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

**Example** The following example modifies the parameters of a connection:  
**connection -modifyset -name MirrorSet -authpolicy -setspecific -username user2 -secret Twelve2Sixteen -o**

## connection -getset

**Lists a specific connection set or all the connection sets**

<b>Description</b>	The <b>naviseccli connection</b> command with the <b>-getset</b> function lists a specific connection set or all the connection sets and displays the contents of each set.
<b>Syntax</b>	<b>connection -getset</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a> ) as follows:  <b>connection -getset [-name <i>name</i>]</b>
<b>Prerequisites</b>	You must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<b>-name <i>name</i></b>  Specifies a unique identifier for the connection set.
<b>Output</b>	<pre> 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.</pre>

Sample output if the storage 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
IP Address of Target: 172.20.4.147
Listening IP port of Target: 3260
SP Ports: B-3v1
Description: Path to tryzub B-1, generated for
MirrorView
```

**Example** The following example lists a specific connection set:  
**connection -getset -name MirrorSet**

## connection -addpath

**Adds a connection path to a specific connection set**

**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** is used with **naviseccli** (described on [page 41](#)) as follows:

```
connection -addpath -name name -target address[:port] [-sp a | b
-portid portid [-vportid vportid]] [comment text] [-o]
```

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

**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 storage system.

**-sp** *a | b*

Specifies the storage processor on the initiating storage system used for the session.

**-portid** *portid*

Specifies the front-end port number on the storage processor of the initiating storage 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.

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

**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
```

## connection -delpath

**Deletes a connection path from a connection set**

**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** is used with **naviseccli** (described on [page 41](#)) as follows:

**connection -delpath -name** *name* **-target** *address[:port]* [**-o**]

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

**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 storage system.

**-o**

Standard override option. Without this option, the user is prompted to confirm intent.

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

**Example** The following example deletes a connection path from a connection set:

**connection -delpath -name MirrorSet -target 123.231.1.1 -o**

## connection -modifypath

**Modifies the parameters of a connection path**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
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]
```

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

**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 storage system.
- sp a | b**  
Specifies the storage processor on the initiating storage system used for the session.
- portid** *portid*  
Specifies the front-end port number on the storage processor of the initiating storage 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.

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

**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
```



## connection -verifypath

Tests the parameters of a connection path

**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.

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

**Syntax** **connection -verifypath** is used with **naviseccli** (described on [page 41](#)) as follows:

**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 storage system.
- sp** *a | b*  
Specifies the storage processor on the initiating storage system used for the session.
- portid** *portid*  
Specifies the front-end port number on the storage processor of the initiating storage 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.

**Ouput** Test/verify command for target {ipaddr[:portnum]} from {fe port} completed successfully.

**Example** The following example tests a connection path from a connection set:  
**Connection -verifypath -name MirrorSet -target 23.231.1.1 -o**

## isns -addserver

**Adds an iSNS server to a storage system's iSNS service configuration**

**Description** The `naviseccli isns` command with the `-addserver` function adds an iSNS server to a storage system's iSNS service configuration.

When you add a server, if it is the first server you add to a storage system's iSNS service configuration, it is designated as the primary server. If a server exists in the storage 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` is used with `naviseccli` (described on [page 41](#)) as follows:

`isns -addserver IPAddr [-o]`

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

**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. Example 111.222.33.44.

---

**-o**

Executes the command without prompting for confirmation.

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

**Example** The following example adds the specified iSNS server, to the storage system's iSNS service configuration:

```
naviseccli -address ss1_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
```

## isns -deleteserver

**Deletes an iSNS server from a storage system's iSNS service configuration**

**Description** The `naviseccli isns` command with the `-deleteserver` function, deletes an iSNS server from a storage system's iSNS service configuration.

---

**Note:** If the IP address of the server you specify to delete is not configured, you will receive an error message.

If you delete a primary server, the next server listed in the storage system's iSNS service configuration becomes the primary.

---

**Note:** The primary server is the server the iSNS service actively communicates with.

**Syntax** `isns -deleteserver` is used with `naviseccli` (described on [page 41](#)) as follows:

`isns -deleteserver IPAddr [-o]`

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

**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. Example 111.222.33.44.

**-o**

Executes the command without prompting for confirmation.

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

**Example** The following example deletes the specified iSNS server from the storage system's iSNS configuration:

```
naviseccli -address ss1_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
```

## isns -listserver

**Displays all iSNS servers that are configured in the storage system's iSNS service**

**Description** The **naviseccli isns** command with the **-listserver** function, lists all servers that are configured in a storage system's iSNS service.

If there are multiple servers configured in the storage system's iSNS service, the **-listserver** operation will list multiple IP addresses and server types. You can only have 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** is used with **naviseccli** (described on [page 41](#)) as follows:

**isns -listserver**

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

**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.

**Example** The following example lists the primary and backup iSNS servers for the storage system:

```
naviseccli -address ss1_spa isns -listserver
```

## isns -setprimary

Sets a configured iSNS server, as the primary server for a storage system's iSNS service configuration

### Description

The `naviseccli isns` command with the `-setprimary` function, establishes a configured iSNS server as the primary server in a storage system's iSNS service configuration. The primary server is the server the iSNS service actively communicates with. You can only have 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.

You use this command when you want to change the primary server in a storage system's iSNS service configuration. You can also use this command to re-establish a primary server after a failover. When a failure occurs which 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
isns -setprimary IPAddr [-o]
```

where:

### Prerequisites

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

### 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. Example 111.222.33.44.



**-o**

Executes the command without prompting for confirmation.

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

**Example** The following example establishes the specified iSNS server as the primary server:

```
naviseccli -address ss1_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
```



---

This chapter explains the Navisphere Manager and Navisphere Express CLI 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.

---

**Note:** This feature is supported only on AX4-5 series and CX series storage systems. It is not supported on the AX series storage systems.

Major sections in the chapter are:

- ◆ [migrate -start](#)..... 420
- ◆ [migrate -cancel](#) ..... 423
- ◆ [migrate -modify](#)..... 424
- ◆ [migrate -list](#) ..... 425

## migrate -start

### Starts a LUN migration session

#### 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 Powerlink.

---

#### Syntax

**migrate -start** is used with **naviseccli** (described on [page 41](#)) as follows:

There are two distinct versions of this command, one for Navisphere Manager 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 Navisphere Manager and requires you to bind the destination LUN before running the command.

---

**Note:** The following command syntax works only with Navisphere Manager.

---

**migrate -start -source** *LUN ID | WWN* **-dest** *LUN ID | WWN* [**-rate** *low | medium | high* | **asap** | *value*]

---

**Note:** The following command syntax works only with Navisphere Express.

---

**naviseccli migrate -start -source** *LUN ID | WWN* **-rg** *ID* [**-o**]

#### Prerequisites

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

**Options** The command switches for Navisphere Manager and Navisphere Express are given below:

The following command switches are supported with Navisphere Manager:

**-source** *LUN ID* | *WWN*

Specifies the source LUN for the migration. You can specify the LUN ID or World Wide Name (WWN).

**-dest** *LUN ID* | *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 storage system is in production, as the normal host I/O may be interrupted. Use **asap** only when the storage system is offline (free of any host-level I/O).

---

The following command switches are supported with Navisphere Express and are mandatory:

**-source** *LUN ID* | *WWN*

Specifies the source LUN for the migration. You can specify the LUN ID or World Wide Name (WWN).

**-rg** *ID*

Specifies the RAID group ID for the destination LUN to be created.

**-o**

Executes the command without prompting for confirmation.

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

**Example Navisphere Manager:**

```
naviseccli migrate -start -source 6 -dest 7 -rate low
```

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.

**Navisphere Express:**

```
naviseccli migrate -start -source 1 -rg 2 -o
```

This command starts a LUN migration session where the source LUN's ID is 1, and the RAID group ID for the destination LUN is 2.

## migrate -cancel

### Cancels a LUN migration session

**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.

For AX4-5 series storage systems, the **migrate** command with the **-cancel** switch is supported only for halted or faulted migrations.

---

**Note:** You can cancel only one migration at a time.

---

**Syntax** **migrate -cancel** is used with **naviseccli** (described on [page 41](#)) as follows:

**migrate -cancel -source** *LUN ID* | *WWN* [-**o**]

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

**Options** **-source** *LUN ID* | *WWN*  
Specifies the source LUN for migration as the LUN ID or World Wide Name (WWN).

**-o**  
Overrides; does not prompt for confirmation.

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

**Example** **naviseccli migrate -cancel -source 6**

This command cancels a LUN migration session whose source LUN's ID is 6.

## migrate -modify

### Modifies a LUN migration session

**Description** The `naviseccli migrate` command with the `-modify` switch modifies a LUN migration session according to the parameters you specify.

**Syntax** `migrate -modify` is used with `naviseccli` (described on [page 41](#)) as follows:

```
migrate -modify -source LUN ID | WWN
[-rate low | medium | high | asap | value] [-o]
```

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

**Options** `-source` *LUN ID | 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 storage system is in production, as the normal host I/O may be interrupted.

Use **asap** only when the storage system is offline (free of any host-level I/O).

---

`-o`

Overrides; does not prompt for confirmation.

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

**Example** `naviseccli migrate -modify -source 6 -rate medium`

This command modifies LUN migration rate to medium.



## migrate -list

### Lists LUN migration sessions and their properties

**Description** The `naviseccli migrate` command with the `-list` switch lists the existing LUN migration sessions and their properties.

**Syntax** `migrate -list` is used with `naviseccli` (described on [page 41](#)) as follows:

```
migrate -list [-source LUN ID | WWN] [-destination] [-rate] [-state] [-percentcomplete] [-timeremaining]
```

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

**Options** `-source LUN ID | 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.

**-destination**

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).

**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)
```

**Example** `naviseccli migrate -list`

This command lists the migration sessions and their properties.

This chapter provides a brief overview of the Navisphere domain and security features, and includes a list of commands you can use to manage storage-system domains and user access to storage systems. *EMC Navisphere Command Line Interface (CLI)* (P/N 069001038) provides a detailed account of Navisphere domain and security features.

Major sections are:

◆ About domains and user accounts .....	429
◆ Domain and security commands overview .....	432
◆ domain -add.....	434
◆ domain -list .....	436
◆ domain -remove .....	437
◆ domain -setmaster.....	439
◆ ntp -list.....	441
◆ ntp -set .....	443
◆ security -adduser.....	446
◆ security -changeuserinfo.....	448
◆ security -list.....	450
◆ security -listrole .....	451
◆ security -rmuser .....	452
◆ security -ldap -addserver .....	453
◆ security -ldap -modifyserver .....	456
◆ security -ldap -removeserver .....	459
◆ security -ldap -listserver .....	460
◆ security -ldap -addrolemapping .....	462
◆ security -ldap -modifyrolemapping .....	464
◆ security -ldap -removerolemapping.....	465
◆ security -ldap -listrolemapping.....	466

- ◆ security -ldap -synchronize ..... 467

---

## About domains and user accounts

A domain is a group of one or more storage systems with Storage Management Server software whose SPs are connected to a network and which have been assigned to the domain by Navisphere CLI or Navisphere Manager. Each domain has a master node (master storage system) that maintains the master copy of the domain data — the storage systems and global user accounts that make up the domain.

Setting up a domain allows a group of storage systems to be monitored and managed using a single login. Even if you plan to use a storage system by itself (manage it separately), we suggest that you create a domain for that system.

---

**Note:** When you set up a user account, you assign a scope of either global, local, or LDAP. Global user accounts apply to all storage systems within a domain. Local user accounts apply to a specific storage system. LDAP uses an external server to authenticate accounts for an entire domain.

---

Navisphere Manager CLI versions 6.24 and later support lightweight directory access protocol (LDAP). This requires an active and operational LDAP environment to which you can connect Navisphere 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 Manager 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 storage system operations) can have one of three roles:

- ◆ Administrator
- ◆ Manager
- ◆ Monitor

**Note:** To issue security CLI commands for user account set-up, modification, or removal you must have administrator privileges. Global administrators can manage both global and local user accounts. Local administrators can only manage local user accounts. [Table 4 on page 430](#) 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 storage systems in the domain) or locally (the user has the role on a specific storage system only). Each global username must be unique in the domain; each local username must be unique within the local management server.

**Note:** Usernames 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.

[Table 4](#) defines the operations a user can perform by role and scope.

**Table 4** Operations that users with different roles can perform

User role	Can view	Can add, modify, or delete
Global administrator	All domain and storage-system settings, and global and local accounts	All domain and storage-system settings, and global and local accounts (cannot delete the last global administrator account)
Local administrator	Local storage-system settings and local user accounts	Local storage-system settings and local user accounts
Global manager	All storage-system settings in domain	All storage-system settings in domain
Local manager	Local management server storage-system information	Local storage-system settings
Global monitor	All storage-system settings and accounts in domain	Nothing
Local monitor	Local storage-system settings and accounts in domain	Nothing

---

## About LDAP/Active Directory

Navisphere 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.

[Table 5](#) provides the list of **security -ldap** commands used to configure the LDAP settings.

## Domain and security commands overview

The **naviseccli** commands in this chapter include the domain and security commands. The following table identifies the commands.

**Note:** The domain commands, **-list**, **-setmaster**, and **-remove**, running on AX4-5 series Navisphere Express are supported with FLARE version 02.23.050.5.7.xx or later.

The respective security commands, **-adduser**, **-changeuserinfo**, **-rmuser**, and **-list** supports adding, changing, removing a global administrator user account information and listing all the users on an AX4-5 series Navisphere Express supported with FLARE version 02.23.050.5.7.xx or later.

Table 5 **naviseccli domain and security commands**

Command	Description
<b>Domain commands</b>	
<b>domain -add</b>	Adds one or more storage systems to a domain.
<b>domain -list</b>	Displays all storage systems in a domain.
<b>domain -remove</b>	Removes a storage system from a domain.
<b>domain -setmaster</b>	Sets a storage system as the master storage system in a domain.
<b>ntp -list</b>	Lists the NTP configuration settings in a domain.
<b>ntp -set</b>	Modifies the NTP configuration in a domain.
<b>Security commands</b>	
<b>security -adduser</b>	Creates a user account.
<b>security -changeuserinfo</b>	Changes the password and/or role of a user account.
<b>security -list</b>	Displays user accounts.
<b>security -listrole</b>	Lists the user's security role.
<b>security -rmuser</b>	Deletes a user account.
<b>Security LDAP commands</b>	
<b>security -ldap -addServer</b>	Creates a new external server.



Command	Description
<b>security -ldap -modifyserver</b>	Modifies the external server login configuration.
<b>security -ldap -removeserver</b>	Removes a server.
<b>security -ldap -listserver</b>	Lists the LDAP external server information.
<b>security -ldap -addrolemapping</b>	Creates a new role mapping.
<b>security -ldap -modifyrolemapping</b>	Modifies the role mapping.
<b>security -ldap -removerolemapping</b>	Removes the role mapping entity.
<b>security -ldap -listrolemapping</b>	Lists the role mapping entity.
<b>security -ldap -synchronize</b>	Synchronizes the accounts.

## domain -add

### Adds one or more storage systems to a domain

#### Description

The **naviseccli domain** command with the **-add** function, adds the storage system you specify to the storage system domain. You can choose to add one or multiple storage systems to the domain. When you add a storage system to the domain, you specify the storage system's IP address.

See **domain -setmaster** if you need to establish a domain. You then can add storage systems to the domain using **domain -add**.

---

**Note:** The **domain -add** command is supported for AX4-5 series Navisphere Express only with FLARE version 02.23.050.5.7.xx or later and not with any prior FLARE versions.

---

#### Syntax

**domain -add** is used with **naviseccli** (described on [page 41](#)) as follows:

**domain -add** *IPAddr* [-o]

#### Prerequisites

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

#### Options

*IPAddr*

The IP address of the storage system you want to add.

---

**Note:** You must identify the storage system by the storage 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. Example 111.222.33.44.

---

**-o**

Executes the command without prompting for confirmation.

#### Output

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

**Example** The following example adds the specified storage systems, to the storage system domain:

```
naviseccli -address ss1_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
```

## domain -list

### Displays all storage systems in a domain

#### Description

The **naviseccli domain** command with the **-list** function, lists the IP address, name, port and secure port information for all storage systems in a storage system domain. If you only want to view information about the master storage system in the domain, you can use the optional **-master** switch.

**Note:** You can only have one master storage system in a domain (see **domain -setmaster**).

The **domain -list** command is supported for AX4-5 series Navisphere Express only with FLARE version 02.23.050.5.7.xx or later and not with any prior FLARE versions.

#### Syntax

**domain -list** is used with **naviseccli** (described on [page 41](#)) as follows:

```
domain -list [-master]
```

#### Prerequisites

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

#### Options

**-master**

Specifies to display only information about the master storage system in the domain.

#### Output

```
Node:          APM000111111111
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
```

#### Example

The following example lists the storage systems in the storage system domain:

```
naviseccli -address 111.222.33.44 domain -list
```

## domain -remove

### Removes a storage system from a domain

**Description** The **naviseccli domain** command with the **-remove** function, removes the storage system you specify from a storage system domain. When you remove a storage system from a domain, you specify the storage system's IP address.

If you want to move a storage 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 storage system from the domain, you destroy the domain. The CLI warns you before you can complete the operation. If the storage system you want to move is the master storage system in the domain, and you want to retain the domain, then you should change the master storage system to another system in the domain before moving the old master (see **domain -setmaster**).

---

**Syntax** **domain -remove** is used with **naviseccli** (described on [page 41](#)) as follows:

**domain -remove** *IPAddr* [-o]

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

**Options** *IPAddr*

The IP address of the storage system you want to remove.

---

**Note:** You must identify the storage system by the storage 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. Example 111.222.33.44.

---

**-o**

Executes the command without prompting for confirmation.

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

**Example** The following example removes the specified storage system, from the storage system domain:

```
naviseccli -address ss1_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
```

## domain -setmaster

Sets a storage system as the master storage system in a domain

**Description** The **naviseccli domain** command with the **-setmaster** function, establishes the storage system you specify as the master storage system in a domain. The master storage system holds the master copy of all global account information in the domain. This information is copied to all member storage systems in the domain. You can only have one master storage system in a domain. You can use **domain -setmaster** to establish a domain or to change the master storage system in an existing domain.

---

**Note:** An AX series storage system can act as the domain master in a single-system domain that contains only an AX series system. An AX series system *cannot* act as domain master in a multiple-system domain. You must establish a CX3 series, CX series, or off-array host to act as the domain master.

---

If you are using **domain -setmaster** to set up a domain, you specify the IP address of a new or existing storage system, not part of an existing domain. This storage system becomes the master and initializes the domain. You can then use the **domain -add** command to add additional storage systems to the domain.

When you use **domain -setmaster** to change the master storage system in an existing domain, the existing master storage system is automatically demoted, and the storage system you specify becomes the new master storage system.

---

**Note:** The **domain -setmaster** command is supported for AX4-5 series Navisphere Express only with FLARE version 02.23.050.5.7.xx or later and not with any prior FLARE versions.

---

**Syntax** **domain -setmaster** is used with **naviseccli** (described on [page 41](#)) as follows:

**domain -setmaster** *IPAddr* [-o]

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

**Options***IPAddr*

The IP address of the storage system you want to set as master.

---

**Note:** You must identify the storage system by the storage 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. Example 111.222.33.44.

---

**-o**

Executes the command without prompting for confirmation.

**Output**

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

**Example**

The following example sets the specified storage system, as the master storage system in a domain:

```
naviseccli -address ss1_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
```



## ntp -list

**Lists the NTP configuration settings in effect in a domain**

**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]`

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

Only a security administrator or an administrator have the access rights.

You will receive an error message if access is denied or an invalid parameter is set.

**Options** `-control`

Displays start/stop and interval information.

`-servers`

Displays configured server addresses and any associated authentication information.

`-all`

Displays all information.

**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"
"1234567812345678"</VALUE>
</PARAMVALUE>
</METHODRESPONSE></SIMPLERSP></MESSAGE></CIM>

```

**Example** The following example lists all the information:

```

naviseccli -user a -password a -scope 0 -address 10.5.1.207 ntp -list -all

```

## ntp -set

### Modifies the NTP configuration in a domain

**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...]]`

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

Only a security administrator or an administrator have access rights.

You will receive an error message if the access is denied or an invalid parameter is set.

**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).

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.

Is required if **serverkey** is specified. If server authentication is not desired, **keyvalue** is not required.

Enclosure key values in quotation marks.

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).

**Output** Varies depending on the switch you use.

**Example** **naviseccli -address ss1\_spa -set -start yes -interval 45 -servers 10.5.1.207 128.222.132.13 -serverkey 1 5 -keyvalue "1234567812345678" "abcdefghabcdefgh"**

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 5 0 -keyvalue "1234567812345678" ""**

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**

Specifies items for **ss1\_spa** choosing not to use server authentication.

**naviseccli -address ss1\_spa -set -start no**  
Stops (disables) NTP operations for **ss1\_spa**.

**Note:** Does not affect synchronization interval or servers settings.

```
naviseccli -address ss1_spa -set -interval 45 -servers  
10.5.1.207 128.222.132.13 -serverkey 5 0 -keyvalue  
"1234567812345678" ""
```

Specifies the interval and server information for ss1\_spa.

Note: Does not affect start/stop settings.

## security -adduser

### Creates a user account

**Description** The `naviseccli security` command with the `-adduser` function, adds a user account to the storage system you specify. If you create a global account, it replicates to all storage systems in the domain. When you add a user account, you specify the username, password, scope and role.

---

**Note:** When you create the initial global administrator user account, you do not need to provide a username and password at login.

The `security -adduser` command for adding a global administrator user account on AX4-5 series Navisphere Express is supported only with FLARE version 02.23.050.5.7.xx and not with any prior FLARE versions.

---

**Syntax** `security -adduser` is used with `naviseccli` (described on [page 41](#)) as follows:

```
security -adduser -user username -password password -scope  
global | local -role administrator | manager | monitor [-o]
```

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

**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 | manager | monitor`

Specifies the role, administrator, manager, or monitor, you want to apply to the user account.

`-o`

Executes the command without prompting for confirmation.

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

**Example** The following example creates the initial global administrator account on the storage system. You are then prompted to create a domain for the storage system you specify:

```
naviseccli -address ss1_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 storage system is not in a domain. It is highly  
recommended you create a new domain for this system.
```

## security -changeuserinfo

**Changes the password and/or role of a user account**

**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.

The `security -changeuserinfo` command for changing global administrator user account information on AX4-5 series Navisphere Express systems is supported only with FLARE version 02.23.050.5.7xx and not with any prior FLARE versions.

---

**Syntax** `security -changeuserinfo` is used with `naviseccli` (described on [page 41](#)) as follows:

```
security -changeuserinfo -user username -scope global | local
[-newpassword password][-newrole administrator | manager |
monitor] [-o]
```

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

**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 storage 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.

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

**Example** The following example modifies the role of an existing user account:  
**naviseccli -address ss1\_spa security -changeuserinfo -user b -scope local -newrole manager**

```
WARNING: You are about to change user:  b (local)
Proceed?(y/n) y
```

## security -list

### Displays user accounts

#### Description

The **naviseccli security** command with the **-list** function, lists the username, scope and role of user accounts. You can filter your selection using any combination of username, scope and role.

**Note:** The **security -list** command is supported on all Navisphere Manager storage systems.

The **security -list** command on AX4-5 series Navisphere Express systems is supported only with FLARE version 02.23.050.5.7xx and not with any prior FLARE versions.

#### Syntax

**security -list** is used with **naviseccli** (described on [page 41](#)) as follows:

```
security -list [-user username] [-scope global | local] [-role  
administrator | manager | monitor]
```

#### Prerequisites

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

#### 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 **administrator** | **manager** | **monitor****

Identifies users with a specific role for which you want to display information.

#### Output

```
Username:  c
Role:      monitor
Scope:     global
```

#### Example

The following example lists users with the specified scope and role:

```
naviseccli -address ss1_spa security -list -role monitor -scope  
global
```

## security -listrole

**Lists the user's security role**

**Description** The `naviseccli security` command with the `-listrole` function, lists the user's security role.

**Syntax** The switch `-listrole` is used with `naviseccli` (described on [page 41](#)) as follows:

`security -listrole`

**Prerequisites** You must have a user account on the storage system on which you want to execute the command. If no user account matches the query, an access denied error will be displayed.

**Output** `Role: Administrator`

**Example** The following example lists the user's security role:  
`naviseccli -h 10.14.83.44 security -listrole`

## security -rmuser

### Deletes a user account

#### 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.

The **security -rmuser** command for removing a global administrator user account on AX4-5 series Navisphere Express systems is supported only with FLARE version 02.23.050.5.7xx and not with any prior FLARE versions.

#### Syntax

**security -rmuser** is used with **naviseccli** (de-scribed on [page 41](#)) as follows:

```
security -rmuser -user username -scope global | local [-o]
```

#### Prerequisites

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

#### 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.

#### Output

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

#### Example

The following example removes the specified user account:

```
naviseccli -address ss1_spa security -rmuser -user b -scope local
```

```
WARNING: You are about to remove user:  b (local)
```

```
Proceed?(y/n) y
```

## security -ldap -addserver

Creates a new external server

**Description** The `naviseccli security -ldap` command with the `-addserver` subcommand creates a new external server login configuration. Only a Navisphere security administrator or administrator can connect your storage 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 storage system with their LDAP username/password. Only two service connections can be configured at a time.

Storage 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
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] [-o]
```

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

**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 users 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.

**-o**

Executes the command without prompting for confirmation.

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

**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`

## security -ldap -modifyserver

Modifies the external server login configuration

**Description** The `naviseccli security -ldap` command with the `-modifyserver` subcommand modifies the settings that connect your storage 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
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]
```

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

**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 users 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.

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

**Example** `naviseccli -h 10.32.123.229 -user a -password a -scope 0 security -ldap -modifyserver 10.5.4.250 -bindpassword crosstor -servertype LDAP`

## security -ldap -removeserver

Removes a server

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`security -ldap -removeserver IPaddress [-o]`

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

**Options** `IPaddress`

Deletes the specified network IP address (IPv4) of the LDAP or the AD server.

`-o`

Executes the command without prompting for confirmation.

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

**Example** `security -ldap -removeserver 10.5.4.250`

## security -ldap -listserver

### Lists the LDAP external server information

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`security -ldap -listserver [-primaryserver | -alternateserver]`

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

**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.

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

**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=adminstrator,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
```

## security -ldap -addrolemapping

Creates a new role mapping

**Description** The `naviseccli security -ldap` command with the `-addrolemapping` subcommand creates a new role mapping entity for the corresponding external server. You must initialize storage-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 Navisphere roles. Once you have defined the server connection settings, and mapped the user's roles, users can log in to the storage system with their LDAP username/password. The roles are administrator, manager, monitor, security administrator, replication, local replication, and replication/recovery.

**Syntax** `security -ldap -addrolemapping` is used with `naviseccli` (described on [page 41](#)) as follows:

```
security -ldap -addrolemapping IPaddress -name name -type
group | user -role role [-o]
```

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

**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 Navisphere Manager.

**-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 Navisphere Manager.

**-o**

Executes the command without prompting for confirmation.

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

**Example** `security -ldap -addrolemapping 10.5.4.250 -name ldap4ipv4 -type user -role administrator`

## security -ldap -modifyrolemapping

Modifies the role mapping

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

`security -ldap -modifyrolemapping IPaddress  
-name name -type group | user -role role [-o]`

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

**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 Navisphere Manager.

**-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 Navisphere Manager.

**-o**

Executes the command without prompting for confirmation.

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

**Example** `security -ldap -modifyrolemapping 10.5.4.250 -name test -type user  
-role monitor`



## security -ldap -removerolemapping

Removes the role mapping entity

**Description** The `naviseccli security -ldap` command with the `-removerolemapping` subcommand removes the role mapping entity from Navisphere Manager for the corresponding external server. Only the Navisphere 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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
security -ldap -removerolemapping IPaddress  
[-name name -type group | user] | [-all] [-o]
```

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

**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.

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

**Example** `security -ldap -removerolemapping 10.5.4.250 -name test -type user`

## security -ldap -listrolemapping

Lists a role mapping entity

**Description** The `naviseccli security -ldap` command with the `-listrolemapping` subcommand lists the role mapping entity for the corresponding external server. It displays if the role mapping is an security administrator, administrator, manager, or monitor. You can use the `-name` switch to list a particular role mapping.

**Syntax** `security -ldap -listrolemapping` is used with `naviseccli` (described on [page 41](#)) as follows:

`security -ldap -listrolemapping IPaddress`  
`[-name name]`

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

**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.

**Output**

```
Name: Test_1
Type:  user
Role:  administrator
Name: Test_2
Type:  group
Role:  security administrator
```

**Example** `security -ldap -listrolemapping 10.5.4.250`

## security -ldap -synchronize

Synchronizes the accounts

- Description** The `naviseccli security -ldap` command with the `-synchronize` subcommand manually synchronizes the accounts with the current service.
- Syntax** `security -ldap -synchronize` is used with `naviseccli` (described on [page 41](#)) as follows:  
`security -ldap -synchronize [-o]`
- Prerequisites** You must have a user account on the storage system on which you want to execute the command.
- Options** `-o`  
Executes the command without prompting for confirmation.
- Output** None if the command succeeds; status or error information if it fails.



This chapter provides a brief overview of the Virtual Provisioning feature and describes the storage pool, thin LUN, and thin commands that you can use if this feature is enabled on your system.

Major sections in the chapter are:

◆ Virtual Provisioning overview .....	470
◆ storagepool -create .....	473
◆ storagepool -list .....	475
◆ storagepool -destroy .....	478
◆ storagepool -modify .....	479
◆ storagepool -expand .....	481
◆ storagepool -cancelexpand .....	482
◆ thinlun -create .....	483
◆ thinlun -list .....	485
◆ thinlun -modify .....	489
◆ thinlun -destroy .....	490
◆ thin -info .....	491

---

## Virtual Provisioning overview

Virtual Provisioning is an optional feature which when activated provides thin provisioning functions. Thin provisioning is the ability to allocate virtual storage resources to hosts and then allocate physical storage capacity to LUN's on demand rather than at the time of creation. It is possible to assign far more virtual storage capacity than the physical resource backing it up. This eliminates the need to host re-provisioning savings manpower and eliminating downtime while preserving storage resources and increasing disk utilization rates. With Virtual Provisioning disks are grouped into pools in which thin LUNs are created. Thin Pools and LUNs are managed by the system which greatly simplifies management. User data is densely packed and balanced across all available spindles. The systems monitors and alerts on pool space consumption and the user re-provisions the pool not each LUN. Virtual Provisioning is supported only on CX4 series storage systems.

---

### Thin pools

A thin pool is a type of storage pool. It contains a set of disks (maximum depends on storage-system type), all with the same redundancy (RAID 5 and RAID 6 only). For more efficient performance, we recommend that all disks in the thin pool have the same capacity. A thin pool shares its user capacity with one or more thin LUNs. You can expand the user capacity of a thin pool by adding disks to it.

---

### Thin provisioning limits

Each release has limits on how many and how large pools can be and how many thin LUNs are supported for each system. Refer to the *Open System Configuration Guide* on Powerlink for complete details of the specific release.

---

### Thin LUNs

A thin LUN is the logical unit of storage that is created in a thin pool. It consumes storage only when data is written to it and competes with other thin LUNs in the pool for the available thin pool storage. The size of the thin LUN that is visible to the host is independent of

the available physical storage in the thin pool. A thin LUN behaves very much like a traditional LUN.

---

## Thin LUN limits

Minimum user capacity	1 GB
Maximum user capacity	14 TB

---

## Current thin provisioning restrictions

- ◆ Thin LUNs cannot be used in a reserved LUN pool or as a clone private LUN (CPL)
- ◆ Thin LUNs cannot be a component in a metaLUN.
- ◆ Thin pools support only RAID 5 and RAID 6 RAID types.
- ◆ Thin pools do not support the hot spare RAID type.

Table 6

### naviseccli thin provisioning commands

Command	Description
<b>Storage pool commands</b>	
<b>storagepool -create</b>	Creates a thin pool.
<b>storagepool -list</b>	Lists the thin pools properties.
<b>storagepool -destroy</b>	Deletes a thin pool.
<b>storagepool -modify</b>	Modifies thin pool properties.
<b>storagepool -expand</b>	Expands the thin pool on the storage system.
<b>storagepool -cancelexpand</b>	Cancels the expansion, if the expansion is failed.
<b>Thinlun commands</b>	
<b>thinlun -create</b>	Creates a thin LUN.
<b>thinlun -list</b>	Lists the thin LUNs and their properties.

<b>thinlun -modify</b>	Modifies the thin LUN properties.
<b>thinlun -destroy</b>	Deletes a thin LUN.
<b>thin -info</b>	Lists the Thin Provisioning feature configuration information.



## storagepool -create

Creates the thin pool

**Description** The `naviseccli 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 storage 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. You must specify the RAID type at the time of creation of thin pool.

**Syntax** `storagepool -create` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli storagepool -create -disks B_E_D List -rtype rtype
-type [thin] [-name name] [-prcntFullThreshold Used Space Prcnt]
[-description description text] [-skiprules]
```

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

**Options** `-disks`

Specifies the list of disks in *Bus X Enclosure Y Disk Z (B\_E\_D)* format. If you do not specify B in the disk format, Bus 0 is assumed.

`-rtype`

Specifies the type of RAID group of the thin pool. The valid RAID groups are `r_5` and `r_6`.

`-type`

Specifies the type of the storage pool. Only thin pool is supported.

`-name`

Sets the unique name for the thin pool. The character length ranges from 1 to 64.

---

**Note:** The characters "\*", "!", and "@" are not allowed while naming a storage pool in LINUX platforms.

---

The character "\*" is not allowed while naming a storage pool for any platforms.

---

**-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%.

**-description**

Sets the description for the thin pool to be created. The character length ranges from 0 to 255.

**-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. The default value is *FALSE*.

**Output** None if the command succeeds; status or error information if initiating the command fails.

**Example** Navisphere Manager

```
naviseccli -h ss1_spa storagepool -type thin -create -disks 0_5 0_6  
0_7 0_8 0_9 -rtype r_5
```

This command creates a thin pool with 5 disks.

## storagepool -list

Lists the thin pool properties and RAID groups

- Description** The `naviseccli storagepool -list` command lists the thin pool properties and RAID groups. For RAID group information, see [“getrg” on page 203](#).
- If you do not specify the options, the system displays all the thin pool properties.
- Syntax** `storagepool -list` is used with `naviseccli` (described on [page 41](#)) as follows:
- ```
naviseccli storagepool -list -name poolname [-id pool id
-type [thin | rg] [-disks] [-luns] [-rtype] [-prcntFullThreshold]
[-description] [-diskType] [-status] [-state] [-currentOp] [-prcntOp]
[-opState] [-rawCap] [-userCap] [-consumedCap] [-availableCap]
[-subscribedCap]
```
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options**
- name *poolname* | -id *pool id***  
Displays the name or ID of the thin pool with the specified name or ID. You cannot use the options `-id` and `-name` together.
  - type *thin* | *rg***  
Specifies the type of thin pool. The **thin** option lists all the thin pool properties. The **rg** option lists the RAID group properties. The **rg** option displays the same output as that of the `getrg` command, see [page 203](#).
  - disks**  
Lists the disks used in each thin pool on the storage system. The disks are listed in `B_E_D` format, with each disk on a new line.
  - luns**  
Lists the IDs of all thin LUNs in the thin pool.
  - rtype**  
Lists the thin pool RAID type.

**-prcntFullThreshold**

Displays the user modifiable used space threshold, beyond which the system generates the alerts.

**-description**

Returns the description of the thin pool.

**-diskType**

Displays the disk type of the components in the pool. If the components have different types, the system displays `Mixed`.

**-status**

Displays the additional descriptive information for the current state of the thin pool.

**-state**

Displays the state of the thin pool.

**-currentOp**

Displays the current operation.

**-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.

**-opState**

Displays the current operation state. The system displays `N/A`, if the current operation state cannot be determined.

**-rawCap**

Displays the raw capacity of the thin pools.

**-userCap**

Displays the user-accessible capacity of the thin pool.

**-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 meta-data (mapping info).

**-availableCap**

Displays the free shared capacity (in blocks and GB).

**-subscribedCap**

Displays the host-visible capacity of the thin pool.

**Output**

```
Thin Pool Name: Thin Pool 0
Thin 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 Percent Completed: 100
Raw Capacity (Blocks): 1401733120
Raw Capacity (GBs): 668.398
User Capacity (Blocks): 1111558400
User Capacity (GBs): 530.032
Consumed Capacity (Blocks): 50334720
Consumed Capacity (GBs): 24.001
Available Capacity (Blocks): 1061223680
Available Capacity (GBs): 506.031
Total Subscribed Capacity (Blocks): 106665345021
Total Subscribed Capacity (GBs): 50862.000
Percent Subscribed: 9596.018
Oversubscribed by (Blocks): 105553786621
Oversubscribed by (GBs): 50331.968
Disks:
Bus 0 Enclosure 0 Disk 6
Bus 0 Enclosure 0 Disk 8
Bus 0 Enclosure 0 Disk 5
Bus 0 Enclosure 0 Disk 7
Bus 0 Enclosure 0 Disk 9
LUNs: 302,3,5,300,9,12,11,2,8,10,301
```

## storagepool -destroy

**Destroys the thin pool**

**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 storage system.

**Syntax** `storagepool -destroy` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli storagepool -destroy -id pool id | -name poolname
-type [thin] [-o]
```

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

**Options** `-id pool id | -name poolname`

Specifies the thin pool ID or name to be destroyed. You cannot use the options `-id` and `-name` together.

`-type`

Specifies the type of the storage pool. Only thin pool is supported.

`-o`

Executes the command without prompting for confirmation.

**Output** None if the command succeeds; status or error information if initiating the command fails.

## storagepool -modify

**Modifies the thin pool**

**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` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli storagepool -modify -id pool id | -name pool name  
-type [thin] [-newName newname] [-description new description]  
[-prcntFullThreshold new user capacity threshold] [-o]
```

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

**Options** `-id pool id | -name poolname`

Specifies the thin pool ID or name to be modified. You cannot use the options `-id` and `-name` together.

`-type`

Specifies the type of the storage pool. Only thin pool is supported.

`-newName`

Sets the unique name for the thin pool. The character length ranges from 1 to 64.

---

**Note:** The characters "\*", "!", and "@" are not allowed while naming a storage pool in LINUX platforms.

The character "\*" is not allowed while naming a storage pool for any platforms.

---

`-description`

Sets the new description to the thin pool. The character length ranges from 0 to 255.

**-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%.

**-o**

Executes the command without prompting for confirmation.

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



## storagepool -expand

Expands thin pool capacity

**Description** The `naviseccli storagepool -expand` command lets you expand the physical capacity of the thin pool on the storage system.

**Syntax** `storagepool -expand` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli storagepool -expand -id pool id | -name pool name  
-disks Disk List in B_E_D Format -type [thin] [-skiprules]
```

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

**Options** `-id pool id | -name poolname`

Specifies the thin pool ID or name to be expanded. You cannot use the options `-id` and `-name` together.

`-disks`

Specifies the list of disks in `B_E_D` format. If you do not specify `B` in the disk format, `Bus 0` is assumed.

`-type`

Specifies the type of the storage pool. Only thin pool is supported.

`-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`.

**Output** None if the command succeeds; status or error information if initiating the command fails.

## storagepool -cancelexpand

Cancels a failed expansion

- 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` is used with `naviseccli` (described on [page 41](#)) as follows:  
`naviseccli storagepool -cancelexpand -id pool id | -name pool name -type [thin]`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options** `-id pool id | -name poolname`  
Specifies the thin pool ID or name. You cannot use the options `-id` and `-name` together.  
`-type`  
Specifies the type of the storage pool. Only thin pool is supported.
- Output** None if the command succeeds; status or error information if initiating the command fails.

## thinlun -create

### Creates thin LUNs

**Description** The **naviseccli thinlun -create** command lets you create thin LUNs and add those thin LUNs to the thin pool. A thin LUN appears as a certain size to a host but does not consume the actual blocks until a host writes to it. A thin LUN's storage capacity increases by using the thin pool when needed.

**Syntax** **thinlun -create** is used with **naviseccli** (described on [page 41](#)) as follows:

```
naviseccli thinlun -create -poolId Thin Pool ID | -poolName Thin Pool Name -sp SP-ID -capacity capacity [-aa auto-assignment] [-sq size-qualifier] [-offset logical-block-address] [-name LUN name] [-l LUN number]
```

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

**Options** **-poolId** *Thin Pool ID* | **-poolName** *Thin Pool Name*

Specifies the thin pool ID or name to which the new LUN should be added. You cannot use these options together.

**-sp** *SP-ID*

Specifies the default SP to which the new thin LUN must belong. The valid values are A and B.

**-capacity** *capacity*

Specifies the additional storage capacity you can add to the thin LUN. It defines the host visible size of the thin LUN. The size qualifier for this option is the **-sq** switch.

**-aa** *auto-assignment*

Enables the storage system to automatically assign the thin LUN to an SP. The status of auto assignment is:

1=Enable

0=Disable

**-sq** *size-qualifier*

Specifies the size qualifier for the capacity of the thin LUN component. The valid values are blocks, GB, MB, and so on.

**-offset** *logical-block-address*

Sets the LUN starting logical block address (LBA) to begin at an offset location on the RAID group to which the LUN should belong.

**-name** *LUN name*

Sets the unique LUN name for the thin LUNs. The character length ranges from 1 to 64.

**-l** *LUN number*

Specifies the LUN number. The lower limit is 0 and the upper limit is the number of LUNs supported on the specific platform.

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

## thinlun -list

### Lists the thin LUN properties

**Description** The `naviseccli thinlun -list` command lists the thin LUN properties. If you do not specify the parameters, the system lists all the values for all the LUNs.

**Syntax** `thinlun -list` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli thinlun -list [-l LUN number | -name LUN name]
[-aa auto-assignment] [-alOwner] [-at] [-userCapacity] [-uid]
[-consumedCapacity] [-default] [-drivetype] [-offset] [-rtype]
[-snapshot] [-brwsp] [-lunbusytickssp] [-lunidletickssp] [-itp] [-slct]
[-rwrsp] [-state] [-status] [-sor] [-nzcra] [-owner] [-etp] [-poolName]
```

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

**Options** `-l LUN number`

Specifies the LUN number. The lower limit is 0 and the upper limit is the number of LUNs supported on the platform.

`-name LUN name`

Specifies the LUN name. The character length ranges from 1 to 64.

**Note:** You cannot specify both the parameters, `-l LUN number` and `-name LUN name`, at the same time.

`-aa auto-assignment`

Enables the storage system to automatically assign the thin LUN to an SP. The status of auto assignment is Enabled and Disabled.

`-alOwner`

Specifies the allocation owner.

`-at`

Specifies the auto-trespass property.

`-userCapacity`

Specifies the user capacity (Blocks and GBs).

**-uid**

Returns the thin LUN unique ID.

**-consumedCapacity**

Specifies the consumed capacity of the thin LUN.

**-default**

Specifies the default owner of the thin LUN.

**-drivetype**

Displays the drive type of the thin LUN (thin pool). The system displays `Mixed` for more than one drive type in the thin pool.

**-offset**

Returns the alignment offset assigned, when the LUN was bound.

**-rtype**

Specifies the RAID group type of the thin LUN.

**-snapshot**

Returns the number of SnapView snapshots.

**-brwsp**

Returns the number of host blocks read and written in an SP A and SP B. This option works only when Statistics Logging is enabled.

**-lunbusytickssp**

Returns the duration for which the disks are busy on the thin LUN in SP A and SP B.

**-lunidletickssp**

Returns the duration for which the disks are idle on the thin LUN in SP A and SP B.

**-itp**

Returns the number of implicit trespasses for SP A and SP B.

**-slct**

Returns the statistics logging current time.

**-rwrsp**

Returns the number of host read and write requests in an SP A and SP B.

**-state**

Displays the state of the thin LUN.

**-status**

Displays the status of the thin LUN.

**-sor**

Returns the sum of outstanding requests in SP A and SP B.

**-nzcra**

Returns the nonzero request count arrivals in the SP A and SP B.

**-owner**

Returns the name of the SP that currently owns the LUN. If none of the SPs owns the LUN, the system displays Unknown.

**-etp**

Returns the number of external trespasses for SP A and SP B.

**-poolName**

Displays the pool name of the pool to which this thin lun belongs.

**Output**

```
LOGICAL UNIT NUMBER 0
Name: Thin LUN 0
UID: 60:06:01:60:3F:A0:1E:00:AC:56:76:54:AB:5C:DD:11
Thin Pool Name: Thin Pool 0
Current Owner: SP A
Default Owner: SP A
Allocation Owner: SP A
User Capacity (Blocks): 10485760
User Capacity (GBs): 5.000
Consumed Capacity (Blocks): 4194560
Consumed Capacity (GBs): 2.000
Raid Type: r_5
Disk Type: Fibre Channel
Offset: 0
Auto-Assign Enabled: DISABLED
Auto-Trespass Enabled: DISABLED
Current State: Ready
Recovery State: Invalid
Recovery Data State: Invalid
Status: OK(0x0)
```

```
Is Faulted: false
Is Transitioning: false
Current Operation: None
Current Operation State: None
Current Operation Status: 0
Current Operation Percent Completed: 0
Statistics Logging Current Time: 07/28/08 09:45:14
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
Snapshot LUNs (if any): N/A
```



## thinlun -modify

### Modifies the thin LUN properties

**Description** The `naviseccli thinlun -modify` command lets you modify the thin LUN properties. You must specify the properties of the thin LUN that has to be modified. You should specify a minimum of one property to modify.

**Syntax** `thinlun -modify` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli thinlun -modify -l LUN number [-name LUN name]  
[-aa auto-assignment] [-sp SP-ID]
```

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

**Options** `-l` *LUN number*

Specifies the LUN number. The lower limit is 0 and the upper limit is the number of LUNs supported on the platform.

`-name` *new LUN name*

Sets the new LUN name. The character length ranges from 1 to 64.

`-aa` *auto-assignment*

Enables the storage system to automatically assign the thin LUN to an SP. The status of auto assignment is:

1=Enable

0=Disable

`-sp` *SP-ID*

Specifies the default SP to which the new thin LUN should belong. The valid values are A and B. If you change the default value, a prompt appears warning about the performance.

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

## thinlun -destroy

**Destroys thin LUNs**

**Description** The `naviseccli thinlun -destroy` command lets you destroy a thin LUN. Only the Administrator and the Manager have permission to destroy a thin LUN.

**Syntax** `thinlun -destroy` is used with `naviseccli` (described on [page 41](#)) as follows:

`naviseccli thinlun -destroy -l LUN number [-o]`

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

**Options** `-l LUN number`

Specifies the LUN number. The lower limit is 0 and the upper limit is the number of LUNs supported on the platform.

`-o`

Executes the command without prompting for confirmation.

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

## thin -info

### Returns information on the thin provisioning

- Description** The `naviseccli thin -info` command lets you (Administrator or Manager) get information about the thin provisioning feature. If you do not specify the options, all the values are listed.
- Syntax** `thin -info` is used with `naviseccli` (described on [page 41](#)) as follows:  
`naviseccli thin -info [-isSupported] [-maxThinPools] [-maxThinLUNs] [-numPools] [-numThinLUNs] [-minThinLUNSize] [-maxThinLUNSize] [-availableDisks]`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options**
- isSupported**  
Displays whether the virtual provisioning is supported on the specified storage system.
  - maxThinPools**  
Displays the maximum number of thin pools that can be created on the storage system.
  - maxThinLUNs**  
Displays the maximum number of thin LUNs that can be created on the storage system.
  - numPools**  
Displays the number of thin pools created on the storage system.
  - numThinLUNs**  
Displays the number of thin LUNs already created on the storage system.
  - minThinLUNSize**  
Shows the minimum size supported on the thin LUN.
  - maxThinLUNSize**  
Shows the maximum size supported on the thin LUN.

**-availableDisks**

Displays list of disks that are available for virtual provisioning.

**Output**

```
Is Supported: true
Driver Version: 0
Driver Library Name: K10MLUAdmin
Driver State: Ready
Max. Thin Pools: 120
Max. Thin LUNs: 4094
Number of Thin Pools: 1
Number of Thin LUNs: 0
Min. Thin LUN Size(Blocks): 2097152
Min. Thin LUN Size(GBs): 1.000
Max. Thin LUN Size(Blocks): 30064771072
Max. Thin LUN Size(GBs): 14336.000
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
```

This chapter provides a brief overview of the Navisphere Event Monitor and includes a list of commands you can use to configure and manage Event Monitor.

Major sections in the chapter are:

|                                   |     |
|-----------------------------------|-----|
| ◆ Event Monitoring overview ..... | 494 |
| ◆ -portal -create.....            | 496 |
| ◆ -portal -list.....              | 497 |
| ◆ -portal -migrate .....          | 498 |
| ◆ -portal -destroy .....          | 499 |
| ◆ -template -list.....            | 500 |
| ◆ -template -import.....          | 501 |
| ◆ -template -export.....          | 502 |
| ◆ -template -destroy .....        | 503 |
| ◆ -template -findconflict .....   | 504 |
| ◆ -template -swap .....           | 505 |
| ◆ -monitor -addsystem .....       | 506 |
| ◆ -monitor -removesystem.....     | 507 |
| ◆ -monitor -listmapping .....     | 508 |
| ◆ -monitor -applytemplate.....    | 509 |
| ◆ -monitor -stoptemplate .....    | 510 |
| ◆ -monitor -getlog.....           | 511 |
| ◆ -monitor -enableresponse .....  | 512 |
| ◆ -monitor -disableresponse ..... | 513 |
| ◆ -monitor -getlogsize.....       | 514 |
| ◆ -monitor -setlogsize .....      | 515 |
| ◆ -monitor -inserttestevent.....  | 516 |
| ◆ -monitor -responsetest.....     | 517 |
| ◆ -monitor -reloadconfig .....    | 520 |

## 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.

[Table 7](#) lists the **event monitor** commands.

**Table 7** navisecli Event Monitor commands

| Command                        | Description                                    |
|--------------------------------|------------------------------------------------|
| <b>-portal -create</b>         | Adds a host agent to the portal system.        |
| <b>-portal -list</b>           | Lists the portal configurations.               |
| <b>-portal -migrate</b>        | Migrates the portals.                          |
| <b>-portal -destroy</b>        | Deletes the host agent from the portal system. |
| <b>-template -list</b>         | Lists the templates.                           |
| <b>-template -import</b>       | Imports the template files.                    |
| <b>-template -export</b>       | Exports the templates.                         |
| <b>-template -destroy</b>      | Deletes the templates.                         |
| <b>-template -findconflict</b> | Finds and resolves template conflicts.         |
| <b>-template -swap</b>         | Swaps two templates.                           |
| <b>-monitor -addsystem</b>     | Adds the storage system to be monitored.       |
| <b>-monitor -removesystem</b>  | Removes the storage system from the monitor.   |
| <b>-monitor -listmapping</b>   | Lists the storage template mappings.           |
| <b>-monitor -applytemplate</b> | Applies the template to the monitor.           |
| <b>-monitor -stoptemplate</b>  | Stops applying the template to the monitor.    |

| Command                          | Description                                       |
|----------------------------------|---------------------------------------------------|
| <b>-monitor -getlog</b>          | Displays the Event Monitor log.                   |
| <b>-monitor -enableresponse</b>  | Enables the response from the specified monitor.  |
| <b>-monitor -disableresponse</b> | Disables the response from the specified monitor. |
| <b>-monitor -getlogsize</b>      | Gets the log size of the monitor.                 |
| <b>-monitor -setlogsize</b>      | Sets the log size of the monitor.                 |
| <b>-monitor -inserttestevent</b> | Inserts an event in the log.                      |
| <b>-monitor -responsetest</b>    | Creates an artificial event.                      |
| <b>-monitor -reloadconfig</b>    | Reloads the template file.                        |

## **-portal -create**

**Adds a host agent to the portal system**

**Description** The **naviseccli -portal -create** command adds the specified host agent to the specified portal system.

**Syntax** **-portal -create** is used with **naviseccli** (described on [page 41](#)) as follows:

```
eventmonitor -portal -create -system portalsystemname -server  
hostagentname
```

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

**Options**

- system *portalsystemname***  
Specifies the portal system.
- server *hostagentname***  
Specifies the host agent.

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

**Example** **eventmonitor -portal -create -system cx380\_123\_227 -server  
10.32.123.205**



---

## **-portal -list**

**Lists the portal configurations**

**Description** The **naviseccli -portal -list** command lists the portal storage system and its hosts.

**Syntax** **eventmonitor -list** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -portal -list**

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

**Output** Portal configuration is listed. Nothing is listed if the portal configuration is not available.

## **-portal -migrate**

### **Migrates the portals**

**Description** The **naviseccli -portal -migrate** command migrates the portal storage system. The host to which the portals are connected is migrated automatically.

**Syntax** **-portal -migrate** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -portal -migrate -source** *source portal system name* **-dest** *destination portal system name*

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

**Options** **-source** *source portal system name*  
Specifies the original portal configuration name.

**-dest** *destination portal system name*

Specifies the destination portal configuration name.

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

**Example** **eventmonitor -portal -migrate -source** *cx380\_123\_223* **-dest** *cx380\_123\_227*

## **-portal -destroy**

**Deletes the host agent from the portal system**

**Description** The `naviseccli -portal -destroy` command deletes the specified host agent from the specified portal system.

**Syntax** `-portal -migrate` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -portal -destroy -system portal system name -server hostname [-o]`

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

**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.

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

**Example** `eventmonitor -portal -destroy -system cx380_123_227 -server 10.32.123.205`

## **-template -list**

**Lists the templates**

**Description** The **naviseccli eventmonitor -template** command with the **-list** function lists all the templates available in the template database of the target storage system.

**Syntax** **-template -list** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -template -list**

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

**Output** The templates in the database are listed. None if the templates are not available.

## **-template -import**

**Imports the template files**

**Description** The `naviseccli eventmonitor -template` command with the `-import` function imports the template file from the file system to template database.

**Syntax** `-template -import` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -template -import -templatefilename filename`

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

**Options** `-templatefilename filename`  
Specifies the template to be imported to the template database of the target storage system.

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

**Example** `eventmonitor -template -import -templatefilename  
D:\Templates\Template_Test.tpl`

## **-template -export**

**Exports the templates**

**Description** The `naviseccli eventmonitor -template` command with the `-export` function exports the template from the template database to file system.

**Syntax** `-template -export` is used with `naviseccli` (described on [page 41](#)) as follows:

```
eventmonitor -template -export -templatename templatename  
-filepath localfilepath
```

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

**Options** `-templatename` *template name*

Specifies the template.

`-filepath` *localfilepath*

Specifies the local file path.

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

**Example** `eventmonitor -template -export -templatename Template_Test`  
`-filepath D:\Templates\`

## **-template -destroy**

**Deletes the templates**

**Description** The `naviseccli eventmonitor -template` command with the `-destroy` function deletes the specified template from the template database on the target storage system.

**Syntax** `-template -destroy` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -template -destroy -templatename templatename [-o]`

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

**Options** `-templatename templatename`  
Deletes the specified template.

`-o`

Executes the command without prompting for confirmation.

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

**Example** `eventmonitor -template -destroy -templatename Template_Test`

## **-template -findconflict**

**Finds and resolves template conflicts**

**Description** The `naviseccli eventmonitor -template` command with the `-findconflict` function lists the conflicting templates. The `-resolve` switch prompts you to resolve them.

**Syntax** `-template -findconflict` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -template -findconflict [-resolve]`

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

**Options** `-resolve`  
Resolves the template conflict.

**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)
```

**Example** `eventmonitor -template -findconflict -resolve`



## **-template -swap**

**Swaps the two templates**

**Description** The `naviseccli eventmonitor -template` command with the `-swap` function swaps the two specified templates. The storage system applied with template A is applied with template B after the swap.

**Syntax** `-template -swap` is used with `naviseccli` (described on [page 41](#)) as follows:

```
eventmonitor -template -swap -templateName template name A  
template name B [-keepResponses]
```

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

**Options** `-templateName` *template name A* *template name B*  
Swaps the specified template names.

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

**Example** `EventMonitor -template -swap -templateName Template_Test_1`  
`Template_Test_2 -keepResponses`

## **-monitor -addsystem**

**Adds the storage system to be monitored**

**Description** The **naviseccli -monitor** command with the **-addsystem** function adds the specified storage system to the monitored storage system list of the specified host agent.

**Syntax** **-monitor -addsystem** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -monitor -addsystem -server *hostname* -system *storage system name***

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

**Options** **-server *hostname***  
Specifies the host agent.

**-system *storage system name***  
Specifies the storage system.

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

**Example** **eventmonitor -monitor -addsystem -server 10.32.123.205 -system cx380\_123\_223**

## **-monitor -removesystem**

**Removes the storage system from the monitor**

**Description** The `naviseccli -monitor` command with the `-removesystem` function removes the specified storage system from the monitored storage system list of the specified host agent.

**Syntax** `-monitor -removesystem` is used with `naviseccli` (described on [page 41](#)) as follows:

```
eventmonitor -monitor -removesystem -server hostname -system  
storage system name
```

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

**Options**

- `-server hostname`  
Specifies the host agent.
- `-system storage system name`  
Specifies the storage system.

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

**Example** `eventmonitor -monitor -removesystem -server 10.32.123.205 -system  
cx380_123_223`

## **-monitor -listmapping**

**Lists the storage template mappings**

**Description** The **naviseccli -monitor** command with the **-listmapping** function lists all the storage template mappings of the specified centralized or distributed monitor.

**Syntax** **-monitor -listmapping** is used with **naviseccli** (described on [page 41](#)) as follows:

```
eventmonitor -monitor -listmapping -server hostname | -system  
storage system name
```

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

**Options** **-server** *hostname*

Lists the mapping information of the centralized monitor.

**-system** *storage system name*

Lists the mapping information of the distributed monitor.

**Output** The mapping information of the specified monitor is listed. None will be listed if the mapping is not available.

**Example** **eventmonitor -monitor -listmapping -server 10.32.123.205**  
**eventmonitor -monitor -listmapping -system cx380\_123\_223**

## **-monitor -applytemplate**

**Applies the template to the monitor**

**Description** The `naviseccli -monitor` command with the `-applytemplate` function applies the template to the specific event monitor.

**Syntax** `-monitor -applytemplate` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -monitor -applytemplate [-server hostname] [-system storage system name] -templatename templatename`

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

**Options** `-server hostname`

Applies the template to the specified centralized monitor as a global template.

`-system storage 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 storage system monitored by the centralized monitor.

`-templatename templatename`

Specifies the template name.

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

**Example** `eventmonitor -monitor -applytemplate -server 10.32.123.205 -system cx380_123_223 -templatename Template_Test`

## **-monitor -stoptemplate**

### **Stops applying the template to the monitor**

**Description** The `naviseccli -monitor` command with the `-stoptemplate` function stops applying the template to the specific event monitor.

**Syntax** `-monitor -stoptemplate` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -monitor -stoptemplate [-server hostname] [-system storage system name] -templatename templatename`

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

**Options** `-server hostname`

Stops applying the template to the specified centralized monitor.

`-system storage 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 storage system and the centralized monitor.

`-templatename templatename`

Specifies the template name.

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

**Example** `eventmonitor -monitor -stoptemplate -server 10.32.123.205 -system cx380_123_223 -templatename Template_Test`

## **-monitor -getlog**

**Displays the event monitor log**

**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** **-monitor -getlog** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -monitor -getlog -server *hostname* [-h] [+N] | [-N]**

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

**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 oldest 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 : clariion
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 : clariion
Domain : cx420_47_52.
```

## **-monitor -enableresponse**

**Enables the response from the specified monitor**

**Description** The `naviseccli -monitor` command with the `-enableresponse` function enables the response from the specified event monitor.

**Syntax** `-monitor -enableresponse` is used with `naviseccli` (described on [page 41](#)) as follows:

`eventmonitor -monitor -enableresponse -server hostname | -system storage system name`

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

**Options** `-server hostname`

Enables the response from the centralized monitor.

`-system storage system name`

Enables the response from the distributed monitor.

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

**Example** `eventmonitor -monitor -enableresponse -system cx380_123_223`  
`eventmonitor -monitor -enableresponse -server 10.32.123.205`



## **-monitor -disableresponse**

**Disables the response from the specified monitor**

**Description** The `naviseccli -monitor` command with the `-disableresponse` function disables the response from the specified event monitor.

**Syntax** `-monitor -disableresponse` is used with `naviseccli` (described on [page 41](#)) as follows:

```
eventmonitor -monitor -disableresponse -server hostname | -system  
storage system name -time duration min [-o]
```

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

**Options** `-server hostname`

Disables the response from the centralized monitor.

`-system storage system name`

Disables the response from the distributed monitor.

`-time`

Specifies the time in minutes.

`-o`

Executes the command without prompting for confirmation.

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

**Example**

```
eventmonitor -monitor -disableresponse -system cx380_123_223  
-time 30
```

```
eventmonitor -monitor -disableresponse -server 10.32.123.205 -time  
30
```

## **-monitor -getlogsize**

**Gets the log size of the monitor**

**Description** The **naviseccli -monitor** command with the **-getlogsize** function gets the log size of the event monitor.

**Syntax** **-monitor -getlogsize** is used with **naviseccli** (described on [page 41](#)) as follows:

**eventmonitor -monitor -getlogsize -server *hostname* | -system *storage system name***

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

**Options** **-server *hostname***

Gets the log size of the centralized monitor.

**-system *storage system name***

Gets the log size of the distributed monitor.

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

**Example**  
**eventmonitor -monitor -getlogsize -server 10.32.123.205**  
**eventmonitor -monitor -getlogsize -system cx380\_123\_223**

## **-monitor -setlogsize**

**Sets the log size of the monitor**

**Description** The `naviseccli -monitor` command with the `-setlogsize` function changes the log size of Event Monitor.

**Syntax** `-monitor -setlogsize` is used with `naviseccli` (described on [page 41](#)) as follows:

```
eventmonitor -monitor -setlogsize -server hostname | -system storage system name -size logsize [-o]
```

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

**Options** `-server` *hostname*

Sets the log size of the centralized monitor.

`-system` *storage 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.

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

**Example** `eventmonitor -monitor -setlogsize -server 10.32.123.205 -size 20000000`

```
eventmonitor -monitor -setlogsize -system cx380_123_223 -size 20000000
```

## **-monitor -inserttestevent**

**Inserts an event in the log**

- Description** The `naviseccli -monitor` command with the `-inserttestevent` function inserts an event into the Event Monitor event log.
- Syntax** `-monitor -inserttestevent` is used with `naviseccli` (described on [page 41](#)) as follows:  
`eventmonitor -monitor -inserttestevent -server hostname`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- 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**

**Description** The `naviseccli -monitor` command with the `-responsetest` function creates an artificial event to test the Navisphere Event Monitor response.

**Syntax** `-monitor -responsetest` is used with `naviseccli` (described on [page 41](#)) as follows:

`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]`

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

**Options** `-server hostname`  
 Specifies the host agent.

- email**
- 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*.
- emailpage**
- 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*.

**-modempage -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 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.

**-snmp -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**

- Description** The `naviseccli -monitor` command with the `-reloadconfig` function reloads the template file.
- Syntax** `-monitor -reloadconfig` is used with `naviseccli` (described on [page 41](#)) as follows:  
`eventmonitor -monitor -reloadconfig -server hostname`
- Prerequisites** For `naviseccli`, you must have a user account on the storage system on which you want to execute the command.
- Options** `-server hostname`  
Specifies the host agent.
- Output** None if the command succeeds; status or error information if it fails.



---

This chapter explains the virtual server commands to configure virtual center/ESX credentials.

Major sections in the chapter are:

|                                         |     |
|-----------------------------------------|-----|
| ◆ Virtual server overview .....         | 522 |
| ◆ server -volmap .....                  | 523 |
| ◆ server -update .....                  | 526 |
| ◆ server -getagent .....                | 527 |
| ◆ server -remoteconfig .....            | 528 |
| ◆ server -register .....                | 531 |
| ◆ hypervisor -addvirtualcenter .....    | 532 |
| ◆ hypervisor -modifyvirtualcenter ..... | 533 |
| ◆ hypervisor -removevirtualcenter ..... | 534 |
| ◆ hypervisor -addesx .....              | 535 |
| ◆ hypervisor -modifyesx .....           | 536 |
| ◆ hypervisor -removeesx .....           | 537 |
| ◆ hypervisor -listesx .....             | 538 |

## Virtual server overview

The virtual server integration feature simplifies the process for assigning a CLARiiON storage 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 ESX server that needs virtualization integration information with a storage system in the domain. To execute host agent commands, the host does not need be attached to the storage system, but network connectivity is required.

Table 8 lists the virtual server commands.

**Table 8** **naviseccli virtual server commands**

| Command                                | Description                                    |
|----------------------------------------|------------------------------------------------|
| <b>server -volmap</b>                  | Lists the volume map information.              |
| <b>server -update</b>                  | Updates the server information.                |
| <b>server -getagent</b>                | Displays the agent information.                |
| <b>server -remoteconfig</b>            | Displays the remote configuration information. |
| <b>server -register</b>                | Displays the agent information on the host.    |
| <b>hypervisor -addvirtualcenter</b>    | Adds virtual center credentials.               |
| <b>hypervisor -modifyvirtualcenter</b> | Modifies virtual center credentials.           |
| <b>hypervisor -removevirtualcenter</b> | Removes virtual center credentials.            |
| <b>hypervisor -addESX</b>              | Adds ESX credentials.                          |
| <b>hypervisor -modifyESX</b>           | Modifies ESX server properties.                |
| <b>hypervisor -removeESX</b>           | Removes ESX credentials.                       |
| <b>hypervisor -listESX</b>             | Lists ESX servers.                             |

## server -volmap

### Lists the volume map information

**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 storage systems attached to the host. If the specified host is a physical host, it displays **lunmapinfo** information for all storage systems attached to the host. Refer to the *EMC Navisphere CLI Reference* (P/N 300-003-628-A10) for information on **lunmapinfo** command.

**Syntax** `server -volmap` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli -h array_a server -volmap -host ipaddress [-local] [-vm]
```

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

**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 storage system A and B.

**Output** `naviseccli server -volmap -host ipaddress`

Sample output (if the host is 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 ESX):**

```
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 ESX):**

```
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
```

**Example** `naviseccli server -volmap -host ipaddress`

## server -update

### Updates the server information

**Description** The **naviseccli server -update** command polls and updates the physical and ESX servers attached to storage system. It also displays the update status.

**Syntax** **server -update** is used with **naviseccli** (described on [page 41](#)) as follows:

```
naviseccli -h sp server -update  
[-host ipaddress [-rescandevices[-o]]]  
[-all [-rescandevices[-o]]] [-status]
```

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

**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 storage 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.

**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
```

**Example** **naviseccli server -update -status**

## server -getagent

**Gets the agent information**

**Description** The `naviseccli server -getagent` command lists the agent information on the host. The command displays the same output as the classic CLI host `getagent` command.

**Syntax** `server -getagent` is used with `naviseccli` (described on [page 41](#)) as follows:

`naviseccli -h sp server -getagent -host ipaddress`

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

**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**

**Description** The `naviseccli server -remoteconfig` command displays and modifies the host agent configuration information. The command scans all the devices in the storage system on a remote host.

**Syntax** `server -remoteconfig` is used with `naviseccli` (described on [page 41](#)) as follows:

`naviseccli -h sp server -remoteconfig -host ipaddress`

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

**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, storage system, comments, and the connection type (SCSI, TTY, and LAN).

`-users`

Displays the information of 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, array 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, array name, and an optional description.

**-managelan**

Adds the LAN devices in the managed device list. Every device includes a device name, array 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.

**-unmanagedev**

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 storage system on a remote host.

**Output** Provides output depending on the options used.

---

## server -register

**Gets the agent information on the host**

**Description** The `naviseccli server -register` command gets the agent information on the host. The command displays the same output as the classic CLI host `register` command.

**Syntax** `server -register` is used with `naviseccli` (described on [page 41](#)) as follows:

`naviseccli -h sp server -register -host ipaddress`

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

**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

- Description** The `naviseccli hypervisor -addvirtualcenter` command adds the virtual center credentials. The virtual center credentials are stored in the storage system.
- Syntax** `hypervisor -addvirtualcenter` is used with `naviseccli` (described on [page 41](#)) as follows:  
`naviseccli -h sp server -hypervisor -addvirtualcenter ipaddress -username user [-password password] [-description description] [-o]`
- Prerequisites** You must have a user account on the storage system on which you want to execute the command.
- 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.
  - 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**

**Description** The `naviseccli hypervisor -modifyvirtualcenter` command modifies an existing virtual center server's credentials.

**Syntax** `hypervisor -modifyvirtualcenter` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli -h sp server -hypervisor -modifyvirtualcenter name |  
ipaddress [-username user -password password] [-description  
description] [-o]
```

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

**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**

- Description** The `naviseccli hypervisor -removevirtualcenter` command removes an existing virtual center server's credentials.
- Syntax** `hypervisor -removevirtualcenter` is used with `naviseccli` (described on [page 41](#)) as follows:  
`naviseccli -h sp server -hypervisor -removevirtualcenter name | ipaddress [-o]`
- Prerequisites** You must have a user account on the storage system on which you want to execute the command.
- 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

- 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** `hypervisor -addesx` is used with `naviseccli` (described on [page 41](#)) as follows:
- ```
naviseccli -h sp server -hypervisor -addesx name | ipaddress  
-username user [-password password] [-o]
```
- Prerequisites** You must have a user account on the storage system on which you want to execute the command.
- Options**
- name | ipaddress**  
Specifies the name or the IP address of the ESX server.
  - username *user***  
Specifies the username of the ESX server.
  - password *password***  
Specifies the password of the ESX server.
  - 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**

<b>Description</b>	The <b>naviseccli hypervisor -modifyesx</b> command modifies an existing virtual center server's credentials.
<b>Syntax</b>	<p><b>hypervisor -modifyesx</b> is used with <b>naviseccli</b> (described on <a href="#">page 41</a>) as follows:</p> <pre><b>naviseccli -h sp server -hypervisor -modifyesx name   ipaddress -username user -password password [-o]</b></pre>
<b>Prerequisites</b>	You must have a user account on the storage system on which you want to execute the command.
<b>Options</b>	<p><b>name   ipaddress</b></p> <p>Specifies the name or the IP address of the ESX server to be modified.</p> <p><b>-username user</b></p> <p>Specifies the new username of the ESX server.</p> <p><b>-password password</b></p> <p>Specifies the new password of the ESX server. The options <b>-username</b> and <b>-password</b> have to be specified to change either one of them.</p> <p><b>-o</b></p> <p>Sends a prompt if you specify the incorrect username/password.</p>
<b>Output</b>	None if the command succeeds; status or error information if it fails.



## hypervisor -removeesx

**Removes ESX credentials**

**Description** The `naviseccli hypervisor -removeesx` command removes existing ESX server credentials.

**Syntax** `hypervisor -removeesx` is used with `naviseccli` (described on [page 41](#)) as follows:

```
naviseccli -h sp server -hypervisor -removeesx name | ipaddress [-o]
```

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

**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**

**Description** The **naviseccli hypervisor -listesx** command lists all servers with IP address and description.

**Syntax** **hypervisor -listesx** is used with **naviseccli** (described on [page 41](#)) as follows:

**naviseccli -h sp server -hypervisor -listesx**

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

**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:
```

This appendix lists the error codes that the Secure CLI returns. It also includes SnapView and MirrorView error codes, although these applications are explained in other manuals.

Major topics are:

- ◆ Secure CLI error codes ..... 540
- ◆ Storage group command error codes ..... 551
- ◆ SnapView, MirrorView, and other command error codes ..... 553
- ◆ Feature command error codes ..... 566

## Secure CLI error codes

Error value (decimal)	Error value (hexadecimal)	Description
<b>General feature command errors</b>		
0	0x0000	CLI success.
1	0x0001	Invalid arguments.
19712	0x4D00	Management Server is not available.
19713	0x4D01	Error encountered during command execution.
19714	0x4D02	Invalid command entered.
19715	0x4D03	Command is running. Client should poll for status.
19717	0x4D05	Invalid speed supplied by user.
19718	0x4D06	Speed not supported on this port.
19728	0x4D10	CIM API Instance is NULL.
19920	0x4DD0	Invalid command.
<b>Domain error codes</b>		
21408	0x53A0	Directory invalid command.
21409	0x53A1	Directory feature not available.
21410	0x53A2	Directory execute errors.
<b>SnapView clone error codes</b>		
19264	0x4B40	Trans object creation failed.
19265	0x4B41	Trans object loss exception.
19266	0x4B42	Transaction exception.
19267	0x4B43	Unable to get specified instance
19268	0x4B44	No Lun numbers available.
19269	0x4B45	Generic exception.
19270	0x4B46	Cloneview name already used.

Error value (decimal)	Error value (hexadecimal)	Description
19271	0x4B47	Protected restore not enabled.
19272	0x4B48	Can not disable fast restore on cloneviews.
19273	0x4B49	Source or cloneLUs exceeded max limit.
19274	0x4B4A	Can not change LU to private.
19275	0x4B4B	Can not add clone driver to LU ext attributes.
19276	0x4B4C	Can not remove clone driver from LU ext attributes.
19277	0x4B4D	LU not available.
19279	0x4B4F	CPL metalun not supported.
19264	0x4B40	SP not primary or controlling.
19265	0x4B41	Source LU is rolling back.
19776	0x4D40	Invalid snap clone command.
19777	0x4D41	over ride needed.
19778	0x4D42	Unexpected clone error occurred.
19779	0x4D43	Clone feature object is not available.
19780	0x4D44	Clone enabling software is not installed.
19781	0x4D45	Specified CloneGroup is not available.
19782	0x4D46	Specified Clone is not available.
19783	0x4D47	No available LUs for cloning on the system.
19784	0x4D48	Specified LU is not available for cloning.
19785	0x4D49	The Clone Private LUNs have not been allocated.
19792	0x4D50	The protected restore is not supported.
19793	0x4D51	The optional fast recovery is not supported.
19794	0x4D52	Can not add clone, the Max number of clones are already present on CloneGroup.
19795	0x4D53	Can not add clone, since a clone in the clone group is syncing or reverse syncing

Error value (decimal)	Error value (hexadecimal)	Description
19796	0x4D54	Unable to sync or reverse sync clone before it be fractured.
19797	0x4D55	Unable to reverse sync clone because it is Out-Of-Sync.
19798	0x4D56	Unable to reverse sync clone because a clone in its CloneGroup is either Synchronizing or Reverse Synchronizing.
19799	0x4D57	Unable to sync clone because a clone in its Clone Group is either Reverse Synchronizing or Reverse Out-Of-Sync.
19800	0x4D58	Unable to deallocate CPL while there is still clone group.
19801	0x4D59	Can not perform create, addclone or allocateCPL on a LU while it is migrating.
19808	0x4D60	MetaLUN(s) cannot be used as clone private LUN(s).
19809	0x4D61	The luns being specified do not match the source luns. They have to be same capacity as the source luns.
19810	0x4D62	Can not perform setfeature on a LU while it is private.
20032	0x4E40	NAVIALU MR name already exists.
20033	0x4E41	NAVIALU MR LU is dirty.
20034	0x4E42	NAVIALU generic exception.
20035	0x4E43	NAVIALU transaction exception.
20036	0x4E44	NAVIALU MR Lu is private.
20037	0x4E45	NAVIALU MR consumed by aggregate.
20038	0x4E46	NAVIALU MR other drivers found.
20039	0x4E47	NAVIALU MR LU in expanding state.
20040	0x4E48	NAVIALU MR LU is shutdown state.
20041	0x4E49	NAVIALU MR ALU is max.

Error value (decimal)	Error value (hexadecimal)	Description
20042	0x4E4A	NAVIALU MR ALU is comp max.
20043	0x4E4B	NAVIALU MR comp in ALU max.
<b>SnapView clone error codes</b>		
1898348624	0x71268050	An existing Clone in the CloneView is reverse-out-of-sync or reverse-syncing. A new Clone cannot be added.
1898348625	0x71268051	An existing CloneView has the same name as the specified CloneView.
1898348965	0x712681A5	An error occurred allocating the Frozen Clones COD resources.
1898349113	0x71268239	Clone is already administratively fractured.
1898349116	0x7126823C	Can't start synchronization or reverse-sync operation from the current Clone image condition.
1898349123	0x71268243	Attempt to destroy a CloneView which still has at least one Clone.
1898349127	0x71268247	Attempt to remove a Clone while it is being synchronized or reverse-synced.
1898349135	0x7126824F	A Clone in the CloneView is currently in the reverse-out-of sync state.
1898349137	0x71268251	An existing Clone in the CloneView is reverse-syncing.
<b>NQM error codes</b>		
19856	0x4D90	NPO feature object is not available.
19857	0x4D91	NPO driver is not installed.
19858	0x4D92	Unknown exception on CIMAPI calls.
19859	0x4D93	Specified CIM class object does not exist or can not be found.
19860	0x4D94	NULL return from CIM calls. NPO Object does not exist or method is not supported.
19861	0x4D95	Max class number NPO supports is reached.
19862	0x4D96	Invalid NPO subcommand.

Error value (decimal)	Error value (hexadecimal)	Description
19863	0x4D97	Invalid NPO command line switches.
19864	0x4D98	Invalid NPO command line switches value.
19865	0x4D99	LUNs have to be specified by WWNs or LUN numbers.
19866	0x4D9A	Too few parameters.
19867	0x4D9B	Too many parameters.
19868	0x4D9C	Creating archive.
19869	0x4D9D	Need stop NPO.
<b>MirrorView error codes</b>		
19872	0x4DA0	Invalid Sync Mirror command.
19873	0x4DA1	Invalid Async Mirror command.
19874	0x4DA2	Sync Mirror feature object is not available.
19875	0x4DA3	Async feature object is not available.
19876	0x4DA4	Sync enabling software is not installed.
19877	0x4DA5	Async enabling software is not installed
19878	0x4DA6	Unexpected Sync Mirror error occur.
19879	0x4DA7	Unexpected Async error occur.
19880	0x4DA8	No available LUs for mirroring on the system.
19881	0x4DA9	Specified LU is not available for cloning.
19882	0x4DAA	Specified mirror source is not available.
19883	0x4DAB	Mirrorview requires same lun size on mirrored images.
19884	0x4DAC	LUN is in a storage group.
19885	0x4DAD	multiple error occurs.
19886	0x4DAE	intentlog allocate error.
19887	0x4DAF	intentlog deallocate error.



Error value (decimal)	Error value (hexadecimal)	Description
19888	0x4DB0	disablepath error.
19889	0x4DB1	layered driver can't be added to LU stack.
19890	0x4DB2	Adding a layered driver to LU stack failed.
19891	0x4DB3	Removing a layered driver to LU stack failed.
19892	0x4DB4	Mirror can't be activated.
19893	0x4DB5	The mirror can't be deactivated (not active).
19894	0x4DB6	Error on fracture image.
19895	0x4DB7	Error on sync image.
19896	0x4DB8	Error on promote image.
19897	0x4DB9	Requested image wasn't found/didn't exist.
19898	0x4DBA	Error on remove image.
19899	0x4DBB	Error on specified LU is used by other feature and can not be chosen as mirror LU.
19900	0x4DBC	Can not perform create, add secondary image or allocate WIL on a LU while it is migrating.
19901	0x4DBD	Can not perform setfeature on a LU while it is private.
19902	0x4DBE	Image not valid.
19903	0x4DBF	Setfeature on thin LUN unsupported.
<b>MirrorView/A error codes</b>		
20288	0x4F40	Generic Exception.
20289	0x4F41	Mirror length name exceeds max.
20290	0x4F42	Mirror name already used.
20291	0x4F43	Mirror invalid number of LUs.
20292	0x4F44	Mirror group name length exceeds maximum.
20293	0x4F45	Mirror group name already used.
20294	0x4F46	Can not add driver.

Error value (decimal)	Error value (hexadecimal)	Description
20295	0x4F47	Can not remove driver.
20296	0x4F48	SP not primary or controlling.
20297	0x4F49	Mirror group name invalid characters.
20298	0x4F4A	LU participating in mirror.
20299	0x4F4B	Mirror Poll failed.
<b>MirrorView/S error codes</b>		
20544	0x5040	Generic exception.
20545	0x5041	Mirror poll failed.
20546	0x5042	Mirror name length exceeds max.
20547	0x5043	Mirror name already used.
20548	0x5044	Mirror invalid number of LUs.
20549	0x5045	Mirror group name length exceeds maximum.
20550	0x5046	Mirror group name already used.
20551	0x5047	Mirror luns not ready.
20552	0x5048	Mirror secondary already exists.
20553	0x5049	Mirror secondary same as primary.
20554	0x504A	Mirror primary is rolling back.
20555	0x504B	Mirror group name invalid.
20556	0x504C	Mirror group description invalid.
20557	0x504D	WIL metalun not supported.
<b>MirrorView consistency group error codes</b>		
21312	0x5340	The maximum number of groups is already defined.
21313	0x5341	Unable to locate the specified group.
21314	0x5342	The following group cannot be destroyed because it still has mirror members:
21315	0x5343	Unable to locate the specified mirror.

Error value (decimal)	Error value (hexadecimal)	Description
21316	0x5344	This mirror cannot be added to the group because it does not have exactly 1 secondary image.
21317	0x5345	Either the -remoteuid or -remoteaddress switch must be used because this mirror has more than 1 secondary image.
21318	0x5346	The remoteuid cannot be the same the array on which the group resides.
21319	0x5347	The mirror has no images on the specified remote array.
21320	0x5348	The mirror is already a member of a group, so it cannot be added to this one.
21321	0x5349	The remote image specified does not reside on the same array as the rest of the group's members.
21338	0x535A	Specified mirror is not a member of the group.
21339	0x535B	The specified missing mirror is not a member of the group.
21340	0x535C	Group name is already used.
<b>iSCSI (connection) error codes</b>		
21280	0x5320	Invalid connection command.
21281	0x5321	Connection feature object is not available.
21282	0x5322	Connection enabling software is not installed.
21283	0x5323	Unexpected connection error occurred.
<b>iSNS error codes</b>		
21456	0x53D0	Invalid iSNS command.
21457	0x53D1	iSNS feature object is not available.
21458	0x53D2	iSNS enabling software is not installed.
21459	0x53D3	iSNS unexpected errors occurred.
<b>LUN migration error codes</b>		
21248	0x5300	Lun migration invalid command.

Error value (decimal)	Error value (hexadecimal)	Description
21249	0x5301	Lun migration feature not available.
21250	0x5302	Lun migration feature not supported.
21251	0x5303	Lun migration execute error.
21252	0x5304	Lun migration spec migration not available.
21253	0x5305	Lun migration spec source not available.
21254	0x5306	Lun migration base can not be hot spare.
21255	0x5307	Lun migration spec dest not available.
21256	0x5308	Lunmigration error code dest lu internal.
21257	0x5309	Lunmigration error code source lu internal.
<b>Arrayconfig error codes</b>		
21536	0x5420	CLI utility capture feature not available.
21537	0x5421	CLI-capture execute error.
21538	0x5422	Duplicate feature not available.
21539	0x5423	Duplicate execute error.
21540	0x5424	Invalid command.
25864	0x6508	Fix unavailable.
25865	0x6509	Rule exception.
25866	0x650a	Access level violation.
25867	0x650b	Dependencies not met.
25868	0x650C	Rule run exception.
25869	0x650D	Enumerate instances not supported.
25870	0x650E	Task still running.
25871	0x650F	Task limits reached.
25872	0x6510	Inconsistent packages.
25873	0x6511	Single task limit.

Error value (decimal)	Error value (hexadecimal)	Description
<b>Event Monitor (emconfig) error codes</b>		
22272	0x5700	EMCONFIG code get usage.
22273	0x5701	EMCONFIG code invalid argument.
22274	0x5702	EMCONFIG code too few parameters.
22275	0x5703	EMCONFIG code too many parameters.
22276	0x5704	EMCONFIG code bad first option.
22277	0x5705	EMCONFIG code unknown option.
22278	0x5706	EMCONFIG code can load file.
22279	0x5707	EMCONFIG code inaccessible cmd.
22280	0x5708	EMCONFIG code not supported.
22281	0x5709	EMCONFIG code command not valid.
22282	0x570A	EMCONFIG code EMSA not found.
22283	0x570B	EMCONFIG code EMSA save config failed.
22284	0x570C	EMCONFIG code invalid template name.
22285	0x570D	EMCONFIG code classic subcmd parse success.
22286	0x570E	EMCONFIG code classic subcmd parse Failure.
<b>MetaLUN error codes</b>		
19840	0x4D80	Invalid MetaLUN command.
19841	0x4D81	MetaLUN feature object is not available.
19842	0x4D82	MetaLUN feature is not enabled.
19843	0x4D83	Unexpected error occurred during MetaLUN command execution.
19844	0x4D84	Specified MetaLUN is not available.
19845	0x4D85	Specified base LU is not available.
19846	0x4D86	Specified LU is not available.
19847	0x4D87	Destroy multiple ALUs failed.

Error value (decimal)	Error value (hexadecimal)	Description
19848	0x4D88	Destroy multiple ALUs was only partially successful.
<b>BAD Block reporting (BRT) Error Plugin Errors</b>		
22016	0x5600	BRT code not accessible.
22017	0x5601	Invalid not enough parameters.
22018	0x5602	BRT code invalid command.
22019	0x5603	BRT code invalid date format.
22020	0x5604	BRT code invalid end date.
22021	0x5605	BRT code no bad blocks.
22022	0x5606	BRT code output file could not be created.
22023	0x5607	BRT code invalid input file.
22024	0x5608	BRT code Invalid file could not be opened.
22025	0x5609	BRT code clean failed in LUNs.
22026	0x560a	BRT code error getting events.
22027	0x560b	BRT code incorrect input file.
22028	0x560c	BRT code unknown exception.
22029	0x560d	BRT code success with errors.
22030	0x560e	BRT code clean session timeout.

## Storage group command error codes

Error value	Description
82	This version of the FLARE software does not support Access Logix (SAN or shared storage) systems.
83	The group name entered does not match any storage groups for this storage system.
84	The HBA UID specified is not known by the storage system.
85	The LUN number specified is not a bound LUN number.
86	Fairness not supported.
87	The retyped password does not match (case sensitivity not checked).
88	Server could not be found, or is not running an agent.
89	Invalid switch combination.
90	The UID (unique id) you entered does not match any storage groups for this storage system.
91	This version of FLARE software does not support the warm reboot feature.
92	Setpath cannot be used to map to the special default storage group.
93	Invalid number of disks in RAID group.
94	This version of FLARE software does not support the Dual Simultaneous Access.
95	This name identifies more than one storage group. Please identify the storage group that you wish to perform this operation on by providing its unique identifier(uid).
96	Inaccessible command.
97	New storage-system serial number must be a 12-character or less alphanumeric value.
98	Cannot perform inquiry.
99	Cannot read block data.
100	Cannot write to file.
101	Cannot open pipe.

Error value	Description
102	The host specified is not known by the storage system.
103	Network name size is too big.
104	Network admin not configured.
105	Not a privileged user.
106	Read only access.
176	Invalid contact information.
108	Invalid host description.
109	Invalid auto configuration.
110	Invalid privileged user list.
111	Invalid interval.
112	Invalid baud rate.
113	Invalid log size.
114	Invalid managed device.
115	Storage-system option not available.
116	Invalid megapoll value.
117	This version of agent does not support Remote Agent Configuration.
118	Invalid use of explicit device names value.
119	The valid value of baud rate is 9600 or 19200.
120	This option is not supported.
121	Error while parsing file.
122	All of the specified users are already in the privileged user list.



## SnapView, MirrorView, and other command error codes

For additional codes, see the CLI manual for your software.

Error value	Description
123	Snapshot does not exist.
124	This version of FLARE software does not support SnapView.
125	Invalid SP name.
126	Reserved LUN pool does not exist.
127	Package number invalid.
128	This version of FLARE software does not support nondisruptive software installation (NDU).
129	Cannot open this file.
130	File does not exist.
131	RAID type must be <b>r1</b> (RAID 1), <b>r3</b> (RAID 3), <b>r5</b> (RAID 5), or <b>r1_0</b> (RAID1/0).
132	Multiple subcommands specified. Check syntax.
133	Disk for PSM must be on DPE bus 0.
134	Configuration does not exist.
135	Configuration already exists.
136	Size specified is too small.
137	Configuration does not exist. Run the navicli <b>initializearray</b> command to configure the system.
138	First option must be a subcommand.
139	Cannot create RAID group for PSM (Persistent Storage Manager).
140	Name or UID (unique ID) is required.
141	Invalid name specified.
142	Image UID is required.
143	Name and LUN are required.
144	Storage-system UID (unique ID) and LUN UID are required.

Error value	Description
145	Mirror not found.
146	Image not found.
147	Synchronized rate cannot be changed since input image is primary.
148	Name and UID both specified.
149	Invalid reserved LUN pool.
150	Invalid session.
151	Session does not exist.
152	Session is stopping.
153	Invalid snapshot.
154	Snapshot does not exist.
155	The -o option requires "-all" or "-filename."
156	Path to store files is required.
157	Cannot specify both "-all" and "-filename."
158	Enter file index or "quit."
159	Invalid input.
160	Index is out of range.
161	File not found.
162	Space not available to retrieve file.
163	Specified feature not supported.
164	Feature must be specified.
165	Cannot specify both '-lun' and '-lunuid' .
166	Invalid storage processor name.
167	PSM (Persistent Storage Manager) is not broken.
168	PSM (Persistent Storage Manager) is broken. Cannot list or create PSM.
169	LUN cannot be unbound.
170	Operation not supported on this type of storage system.

Error value	Description
171	Incompatible arguments. Invalid storage-system serial number.
172	Directory not specified.
173	Invalid number of blocks.
174	Number of blocks not specified.
175	Reading of data not supported on this storage system.
176	Invalid snapshot World Wide Name (WWN).
177	Invalid storage-system serial number.
178	Navicli '-f' option required to store data in file.
179	Invalid IP address format.
180	Storage group cannot be shared.
181	Invalid HLU number.
182	Invalid ALU number.
183	Invalid port ID.
184	Remote server cannot be managed.
185	Email response test failed.
186	Emailpage response test failed.
187	Modempage response test failed.
188	SNMP response test failed.
189	Phone home response test failed.
190	Mandatory switch for email/emailpage.
191	Mandatory switch for modempage.
192	Mandatory switch for SNMP.
193	Only one message or file can be specified.
194	Valid dial string contains only digits, parentheses, hyphen.
195	File does not exist or cannot open.
196	Specified user already exists.

Error value	Description
197	The offset switch is not supported for this storage system.
198	Valid COM port number is 1,2,3,or 4.
199	Valid dial command is atd, atDp or atD.
200	Valid message delay contains only ",", " (one or more commas).
202	Target LUN number is missing.
203	Session name is missing.
204	SnapView multiple session feature is not supported.
205	Cannot specify both snapshot name and snapshot ID.
206	Cannot specify both -mode and -simulation.
207	This command is not supported on remote host.
208	Switch -pathname must be specified.
209	Get local server attributes failed.
210	This version of FLARE software does not support Hi5 RAID type.
211	Only one of the switches -snapshotid, -snapshotname, or -lun can be specified.
212	Specified session and snapshot must be based on the same target LUN.
213	Cannot add LUN to reserved LUN pool; the maximum number of reserved LUNs has been allocated.
214	Cannot add LUN to reserved LUN pool; the maximum number of reserved LUNs that can be added is <i>n</i> .
215	The HLU (host LUN) number you specified cannot be found.
216	This command must be issued from the SP to which the LUN will trespass.
217	Invalid bus or enclosure number.
218	Invalid WWN seed.
219	Invalid EMC part number.
220	This RAID group has maximum number of LUNs already.
221	Not enough reserved LUNs in the specified SP's reserved LUN pool.

Error value	Description
222	This LUN cannot be added to the Storage Group since it is participating in a remote mirror.
223	Allocate log must specify LUN with valid owner.
224	This request has been issued through the SP that is not the current owner of the targeted LUN.
226	Invalid NDB password.
227	Insert test event failed.
228	The <b>-addroffset</b> switch is supported only for a non-destructive bind.
229	The <b>-addroffset</b> switch must be supplied for a non-destructive bind.
230	Cannot unbind LUN, PSM LUN that is not double faulted.
231	Cannot unbind LUN because it is in use as an active hot spare.
232	Cannot unbind LUN because a feature of the storage system is using it.
233	Cannot unbind LUN because it is contained within a storage group.
234	Incorrect LUN in storage group.
235	Cannot display LUN WWNs and update LUN map information at the same time.
236	The DMF is not supported.
237	Incorrect parameters while creating DMF.
238	Incorrect Name and UID specified for the DMF.
239	It is required to enter Name or UID for the DMF.
240	The specified DMF is removed, the destination does not exist.
241	Incorrect destination LUN WWN and number specified.
242	DMF name is not unique.
243	Unknown Descriptor for DMF.
244	Incorrect List and Set specified for DMF.
245	Throttle switch for DMF is missing.
246	Throttle switch for DMF is invalid.

Error value	Description
247	Invalid Max value for DMF.
248	Invalid poll value for DMF.
249	Incorrect source LUN, WWN and Number specified.
250	The destination for DMF is less than the source.
251	The peer SP is alive.
252	The DMF specified is neither Source nor destination.
253	The specified descriptor name for DMF is duplicate.
254	Duplicate. LUN ID for DMF.
255	The destination DMF exists.
256	DMF copy length is zero.
257	Source specified for DMF is a private LUN.
258	Destination specified for DMF is a private LUN.
259	DMF Modify destination is smaller than source.
260	DMF source specified is not on peer.
261	DMF destination specified is not on peer.
262	SNAP Rollback is not supported.
263	Flush rate is missing for SNAP.
264	Flush rate for SNAP is invalid.
265	SNAP session is not rolling back.
266	SNAP session is not persistent.
267	SNAP session is rolling back.
268	SNAP session source is rolling back.
269	SNAP session has max.
270	SNAP has invalid session name.
271	Simulation is not supported.
272	SNAP chunk size is not changeable.

Error value	Description
273	Session is rolling back, can not stop.
274	Session is rolling back can not activate.
275	Session is rolling back can not deactivate.
276	Mirrors with same name.
277	Error: <b>-phonehome</b> response test is not supported on the targeted agent. Please use <b>-ipconnecthome</b> or <b>-mdmconnecthome</b> instead.
278	Response using modem from agent on SP is not supported.
279	Error: <b>-ipconnecthome</b> response test failed.
280	Error: <b>-ipconnecthome</b> is not supported.
281	Error: <b>-mdmconnecthome</b> response test failed.
282	Error: <b>-mdmconnecthome</b> is not supported.
283	Error: <b>-connecto</b> and <b>-routerip</b> are mandatory.
284	Error: <b>-modemnum</b> , <b>-connectto</b> , and <b>-method</b> are mandatory.
285	Error: <b>-smtpmailserver</b> , <b>-destaddress</b> , and <b>-serialnumber</b> are mandatory.
286	Error: <b>-emailhome</b> response test failed.
287	Error: invalid modem method. Valid method should be one of <b>\\"Direct Dial\\", \\"UUNET WAN\\", \\"UUNET SYM\\", \\"UUNET EQM\\", \\"UUNET JTO\\", \\"UUNET EQT\\", \\"UUNET GEN\\", \\"UUNET GRC\\"</b> and case sensitive.
288	Cannot add a hot spare to a storage group.
289	<b>Chglun</b> parameter not supported on a non-FLARE LUN.
290	Unable to bind the metaLUN, use CLE command <b>metalun -destroy</b> instead.
291	This version of FLARE software does not support <b>getsniiffer</b> or <b>setsniiffer</b> on metaLUNs.
292	You must issue this command from the SP that owns the LUN on which the verify will be run.
293	Cannot specify both <b>-feature</b> and <b>-featurename</b> .
294	Snap_sp_does_not_own_session

295	Snap_sp_does_not_own_snapshot
296	Snap_inactive_snapshot
297	Snap_invalid_snapshot_name
298	Snap_bad_lun_for_snapshot
299	dmf_invalid_lun_or_slv
300	dmf_invalid_sp_name
301	Dmf_dupl_isc_not_supported
302	Dmf_command_not_supported_on_isc
303	Dmf_cannot_mark_if_marked
304	Dmf_cannot_unmark_if_unmarked
305	Dmf_err_nomark_if_marked
306	Dmf_err_copywholelun_if_marked
307	Dmf_err_isc_srclun_not_on_sp
308	Dmf_err_isc_snap_src
309	Dmf_changesonly_isc_not_supported
310	Dmf_copywholelun_isc_not_supported
311	Dmf_nomark_isc_not_supported
312	Dmf_chgtype_invalid_value
313	Dmf_both_switches_error
314	Dmf_nomark_invalid_value
315	Dmf_isc_not_supported
316	Dmf_invalid_linkbw_value
317	Dmf_invalid_linkbw_value_range
318	Dmf_invalid_latency_value
319	Dmf_invalid_granularity_value
320	Dmf_create_isc_params
321	Dmf_modify_isc_params



322	Dmf_isc_name_change
323	Dmf_non_isc_latency_change
324	Dmf_non_isc_granularity_change
325	Dmf_non_isc_linkbw_change = 325,
326	Reserved_snap_sess_invalid_switch = 326,
327	Invalid switch specified for a reserved snap LUN.
328	Cannot activate reserved snapshot LUN.
329	Cannot deactivate reserved snapshot LUN.
330	Cannot stop a reserved session.
331	Cannot remove a reserved snapshot LUN.
332	Cannot modify a reserved snapshot LUN.
333	Cannot start rollback on a reserved session.
334	Cannot change rollback session as it is reserved.
335	The format of decimal number specified is inappropriate.
336	Precision of the decimal number specified exceeds the maximum precision of <i>n</i> (where <i>n</i> varies depending upon the situation).
338	Cannot add a snapshot to storage group because the snapshot specified is reserved.
344	Use the <b>-nolocal</b> switch if agent version is not at least version 6.5 when setting SPS time.
345	Do not use the <b>-nolocal</b> switch if agent version is version 6.5 or greater when setting SPS time. This storage system does not support any NVRAM card.
346	This storage system does not support any NVRAM card.
347	This storage system does not support any operation on shutting down the array.
348	This storage system does not support any operation on LUN cache dirty condition.
349	Error parsing <b>.lst</b> file.
350	Cannot open <b>.lst</b> file.

Error value	Description
351	Syntax error in .lst file.
352	Snap consistent start not supported.
353	The LUN list contains duplicate LUN number.
354	DMF Invalid Autorestart value.
355	Snap migration LUN for snap session.
356	Snap migration LUN for snapshot
357	Cannot port switches together.
358	Snap no hotspare in addluntocache.
359	Snapshot name too long.
360	<b>managedby</b> not supported.
361	<b>cachecard</b> not supported.
362	Invalid SQ flags on AX.
363	DMF destination does not exist.
364	Invalid disk HA.
365	Invalid disks NONHA.
366	Can not access host.
367	DMF lite src not local.
368	DMF invalid port WWN.
369	DMF modify src multi switches
370	Port removedHBA logged in.
371	Luncache no cache dirty.
372	SG not connected to host.
373	Cannot unbind why double db fault.
374	Luncache not default owner.
375	Snap only persistent sessions.
376	DMF DESTPORTWWN DESTWWN and DESTLUN specified.

Error value	Description
377	Need add audit.
378	No luns in RG.
379	RG does not exist.
380	SIW rules failed.
381	No luns owned by target in RG.
382	No luns owned by target.
383	Remove RG duplicate Raid group number.
384	No sniff hot spare.
385	SIW <b>jre</b> not found.
386	SIW rules security.
387	RAID RebootPeerSP command.
388	Unknown SIW error.
389	Illegal error.
390	SIW rules login failed.
391	Resume not supported.
392	RAID reset and hold command.
393	Invalid parity elements hold value.
394	Invalid disk FW.
395	Snapcopy globalcache not committed.
396	Snapcopy globalcache not supported.
397	Non redundant RAID group.
398	Disk not found.
399	Reserved snap can not add lun private.
400	Reserved snap maximum allowed is exceeded.
401	Invalid RG type flags.
402	RAID type supported.

Error value	Description
403	RAID type not supported.
404	RAID type and disk unmatched.
405	Numbers disk even.
406	Invalid Rg type flags 2.
407	DMF name and descriptorID specified.
408	Snap remove active snapshot.
409	Snap active snapshot.
410	The destination LUN does not exist.
411	The same connection type specified.
412	Invalid storage group name.
413	Invalid connection.
414	Remote subsystem cannot be retrieved.
415	Duplicate switches.
416	Maximum of 240 characters allowed.
417	Values not numeric.
418	Inserting event failed.
419	Invalid speed request.
420	Invalid speed duplex.
421	Must be an IPv6 address in colon hexadecimal format.
422	No SP ports.
423	Duplicate disks.
424	Secure CLI executable not found.
425	Cannot_unbind_thinlun
426	Switch_not_supported_metalun
427	Switch_not_supported_thinlun
428	rg_is_private

429	Eprivate_not_supported
430	Luncache_not_supported_metalun
431	Luncache_not_supported_thinlun
432	Not_support_sniffer_on_thinlun
433	Lun_is_in_use
434	Setfeature_sancopy_on_thinlun_not_supported
435	Netadmin_community_too_short
436	Netadmin_community_too_long

## Feature command error codes

The following tables list decimal and hexadecimal values and descriptions of feature command errors, which includes mapped RAID errors, LUN migration errors, connection and iSNS (iSCSI) errors, and domain and security errors.

### General feature command errors

Error value (decimal)	Error value (hexadecimal)	Description
19713	0x4D01	Error encountered during command execution.
19714	0x4D02	Invalid command entered.
19715	0x4D03	Command is running. Client should poll for status.

## Mapped RAID errors

Error value (decimal)	Error value (hexadecimal)	Description
19840	0x4D80	Invalid metaLUN command.
19841	0x4D81	MetaLUN feature object is not available.
19842	0x4D82	MetaLUN feature is not enabled.
19843	0x4D83	Unexpected error occurred during metaLUN command execution.
19844	0x4D84	Specified metaLUN is not available.
19845	0x4D85	Specified base LU is not available.
19846	0x4D86	Specified LU is not available.
19847	0x4D87	Destroy multiple ALUs failed.
19848	0x4D88	Destroy multiple ALUs was only partially successful.

## LUN migration errors

Error value (decimal)	Error value (hexadecimal)	Description
21248	0x5300	Invalid LUN migration command.
21249	0x5301	LUN migration feature object is not available.
21250	0x5302	LUN migration feature is not enabled.
21251	0x5303	Unexpected error occurred during LUN migration command execution.
21252	0x5304	LUN migration object not available.
21253	0x5305	LUN migration source LU is not available.
21254	0x5306	LUN migration base cannot be a hot spare.
21255	0x5307	LUN migration destination LU is not available.

## Connection and iSNS (iSCSI) errors

Error value (decimal)	Error value (hexadecimal)	Description
21280	0x5320	Invalid connection command.
21281	0x5321	Connection feature object is not available.
21282	0x5322	Connection software is not enabled.
21283	0x5323	Unexpected error occurred during connection command execution.
21456	0x53D0	Invalid isns command.
21457	0x53D1	iSNS feature object is not available.
21458	0x53D2	iSNS enabling software is not installed.
21459	0x53D3	Unexpected error occurred during isns command execution.

## Domain and security errors

Error value (decimal)	Error value (hexadecimal)	Description
21408	0x53A0	Invalid domain command.
21409	0x53A1	Domain feature object is not available.
21410	0x53A2	Unexpected error occurred during domain command execution.
21488	0x53F0	Security error code invalid command.
21489	0x53F1	Security error code feature not ava.
21490	0x53F2	Security error code execute error.



## Thin provisioning errors

**Note:** If parsing errors occur, the CLI displays the parameter and the error will appear in the following format, <Parameter Name>: <Error Message>

Error value (decimal)	Error value (hexadecimal)	Description
28049	0x6D91	Internal error. Rule input message could not be created. Please try again.
28050	0x6D92	The best practices check for the storagepool's creation did not finish successfully. The storagepool will not be created.
28051	0x6D93	Internal error. The reference names for some of the disks in the disks list could not be obtained. Please try again.
28052	0x6D94	This storagepool can't be destroyed because it has luns defined in it. Please destroy the luns and then destroy the storagepool.
28053	0x6D95	The best practices check for the storagepool's expansion did not finish successfully. The storagepool will not be expanded.
28054	0x6D96	Could not set properties:(%s). [Parameter 1: A comma separated list of the user visible names of all the properties that could not be set]
28055	0x6D97	The LUN ID list is empty. At least one LUN must be specified.
28056	0x6D98	The following LUNs could not be destroyed: (%s). Please make sure that all these LUNs exist, are private Flare LUNs and belong to the storagepool specified. [Parameter 1: A comma separated list of lun ids that could not be destroyed successfully]
28057	0x6D99	The RAID Group ID list is empty. At least one RAID Group must be specified.

Error value (decimal)	Error value (hexadecimal)	Description
28058	0x6D9A	The following RAID Groups could not be destroyed:(%s). Please make sure that all these RAID Groups exist, are private RAID Groups and belong to the storagepool specified. [Parameter 1: A comma separated list of RG ids that could not be destroyed successfully]
28059	0x6D9B	This command can be used only for Thin LUNs. LUN (%d) is not a Thin LUN. [Parameter 1: The id of the LUN]
28060	0x6D9C	One of the following messages: 1. Cannot unbind LUN because its a PSM LUN that is not double faulted. 2. Cannot unbind LUN because its in use as an active Hot Spare 3. Cannot unbind LUN because its being used by a feature of the Storage System 4. Cannot unbind LUN because its contained in a Storage Group 5. Cannot unbind LUN because the database drive is double faulted.
28061	0x6D9D	Unable to create default name for the Thin LUN. Thin LUN was not created. Please try creating the Thin LUN with the name parameter specified.
28288	0x6E80	Internal error. Could not retrieve: (%s) [Parameter 1: The user visible name of the object class]
28289	0x6E81	Internal error. Could not retrieve: (%s) [Parameter 1: The user visible name of the object class]
28290	0x6E82	Could not retrieve the specified (%s). The (%s) may not exist. [Parameter 1: the user visible name of the object class] [Parameter 2: Same as parameter 1]
28291	0x6E83	Internal error. Could not retrieve the specified (%s). The (%s) may not exist [Parameter 1: the user visible name of the object class] [Parameter 2: Same as parameter 1]

Error value (decimal)	Error value (hexadecimal)	Description
28293	0x6E85	Could not retrieve bus, enclosure and disk position information from:(%s) because the disk names are not in the expected format. [Parameter 1: The string which caused the error]
28294	0x6E86	The following disk could not be retrieved : (%s). [Parameter 1: the valid disk B_E_D string which was passed as input]
19720	0x4D08	Could not retrieve: (%s) [Parameter 1: The user visible name of the object class]
19721	0x4D09	Could not retrieve the specified (%s). The (%s) may not exist. [Parameter 1: the user visible name of the object class] [Parameter 2: Same as parameter 1]
19722	0x4D0A	Internal error. Could not retrieve the specified (%s). The (%s) may not exist [Parameter 1: the user visible name of the object class] [Parameter 2: Same as parameter 1]
19723	0x4D0B	Internal error. Unknown exception while parsing.
19719	0x4D07	This command uses unsupported options. Using these may cause normal system operation to be disrupted and may lead to a data unavailable/data lost (DU/DL) situation. Are you sure you want to perform this operation?(y/n):
19729	0x4D11	Unmatched quotes found: (%s). [Parameter 1: the string with the unclosed quotes]
19731	0x4D13	Internal error. Unable to finish parsing
19732	0x4D14	Internal error. Factory not initialized at (%s). [Parameter 1: the token at which the factory class change was tried]
19733	0x4D15	Internal error. Option handler not initialized at (%s). [Parameter 1: the token at which the option handler was not found]

Error value (decimal)	Error value (hexadecimal)	Description
19734	0x4D16	Internal error. Parameter handler not initialized at (%s). [Parameter 1: the token at which the parameter handler object was not found]
19735	0x4D17	Unrecognized option: (%s). [Parameter 1: the token which wasn't recognized as an option or a parameter of the previous option]
19736	0x4D18	String is too long.<Usage>
19737	0x4D19	String is too short.<Usage>
19738	0x4D1A	String contains characters that are not allowed.<Usage>
19739	0x4D1B	String %s has invalid pattern.<Usage> [Parameter 1: The string that has the wrong pattern]
19740	0x4D1C	Value (%s) out of range.Valid values are (where each value can be specified only once)*: (%s). The remaining valid values are (%s) [parameter 1: The string which is out of the valid list. Parameter 2: A comma separated list of valid string values Parameter 3: A comma separated list of valid string values that can still be used (in case duplication is not allowed *-->This statement in brackets appears only if duplication is not allowed while inputting the string values]
19741	0x4D1D	String (%s) starts with spaces. [Parameter 1: The string that starts with spaces]
19742	0x4D1E	String (%s) ends with spaces. [Parameter 1: The string that ends with spaces]
19743	0x4D1F	Internal error. Unknown pattern initialization
19744	0x4D20	Value %s is not numeric. [Parameter 1: The string that was supposed to be numeric]

Error value (decimal)	Error value (hexadecimal)	Description
19745	0x4D21	Value %llu is less than the minimum. The minimum value is %llu [parameter 1: the numeric value that's less than the minimum Parameter 2: the minimum acceptable numeric value]
19746	0x4D22	Value %llu is greater than maximum. The maximum value is %llu [Parameter 1: The numeric value that's greater than the maximum acceptable value Parameter 2: the maximum acceptable numeric value]
19749	0x4D25	At least one out of the following options should be specified:(%s). [Parameter 1: A comma separated list of options out of which at least one must be specified]
19750	0x4D26	At most one out of the following options should be specified: (%s). [Parameter 1: A comma separated list of options out of which at most one can be specified]
19751	0x4D27	Exactly one out of the following options must be specified: (%s). [Parameter 1: A comma separated list of options out of which exactly one can be specified]
19752	0x4D28	All of the following options must be specified: (%s). [Parameter 1: A comma separated list of options, all of which must be specified]
19753	0x4D29	Too few parameters. At least %d parameters are required. [Parameter 1: The minimum number of acceptable parameters for an option]
19754	0x4D2A	Too many parameters. At most %d parameters can be given. [Parameter 1: The maximum number of acceptable parameters for an option]
19755	0x4D2B	Internal error. While parsing, unable to save value for: %s [Parameter 1: A comma separated list of values which had to be stored]

## Event Monitor errors

Error value (hexadecimal)	Description
0x5460	Info
0x5461	Execute error
0x5462	Feature not ava
0x5463	No system
0x5464	No portal
0x5465	No host agent
0x5466	Portal config already assigned
0x5467	No portal config
0x5468	Portal config not highest version
0x5469	No centralized monitor
0x546A	No distributed monitor
0x546B	No emsa
0x546C	Illegal template name
0x546D	Template already exist
0x546E	Template not applied
0x546F	System not monitored
0x5470	Cannot get emsa config
0x5471	Cannot update emsa config
0x5472	New te mplate name exist in database
0x5473	New template name exist in config
0x5474	Cannot delete call home template
0x5481	Invalid command
0x5482	Illegal template file
0x5483	Cannot upload file

Error value (hexadecimal)	Description
0x5484	Cannot download file
0x5485	No file path
0x5486	Template file already exist

---

## Virtual Server CLI errors

Error value (hexadecimal)	Description
0x7800	Esx success
0x7801	Esx failure
0x7802	Vc success
0x7803	Vc failure
0x7804	Agent success
0x7805	Agent failure
0x7806	Unknown exception
0x7807	Read virtual cfg from psm failure
0x7808	Poll servers failure
0x7809	Poll servers success
0x780a	Init virtual cfg in psm failure
0x780b	Connection failure not vmware vc
0x780c	Connection failure not vmware esx
0x780d	Poll vmware esx autopush success
0x780e	Poll vmware esx autopush failure
0x7900	Command not supported

Error value (hexadecimal)	Description
0x7901	No esx server connection
0x7902	Server does not exist.
0x7903	Not a Virtual Center.



## Secure CLI Command Coverage

---

---

This appendix provides a list of commands that are supported by Secure CLI on various storage systems. The version number indicates the earliest Navisphere version that supported the command for each system type. Unless specified, the command supported once will be supported on all subsequent releases. The commands that are not supported on specific systems are also listed.

- ◆ [Command coverage.....](#) 578

## Command coverage

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
alpa -get	6.19	6.22	6.28	6.23	N/A	Not supported on iSCSI storage system
alpa -set	6.19	6.22	6.28	N/A	N/A	
arraycommpath	6.19	6.22	6.28	6.23	N/A	
arrayconfig -capture	6.26	6.26	6.28	N/A	N/A	
arrayconfig -duplicate	6.26	6.26	6.28	N/A	N/A	
arrayname	6.19	6.22	6.28	6.23	6.23	
backendbus -analyze	N/A	6.22	6.28	N/A	N/A	
backendbus -get -speeds	N/A	6.22	6.28	N/A	N/A	
backendbus -get -sfpstate	N/A	6.22	6.28	N/A	N/A	
backendbus -get -physical	N/A	6.22	6.28	N/A	N/A	
backendbus -get -all	N/A	6.22	6.28	N/A	N/A	
backendbus -resetspeed	N/A	6.22	6.28	N/A	N/A	
baseuuid	6.19	6.22	6.28	N/A	N/A	
bind	6.19	6.22	6.28	6.23	6.23	
cachecard -initialize	N/A	N/A	N/A	N/A	N/A	
cachecard -list	N/A	N/A	N/A	N/A	N/A	
chglun	6.19	6.22	6.28	6.23	6.23	This command does not support thin LUNs.

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
chgrg	6.19	6.22	6.28	6.23	6.23	This command does not support private RAID groups used in thin pools.
clearlog	6.19	6.22	6.28	6.23	6.23	
clearstats	6.19	6.22	6.28	N/A	N/A	
cmdtime	6.19	6.22	6.28	N/A	N/A	
copytohot spare	6.26	6.26	6.26	N/A	N/A	Not supported on AX4-5 series
createrg	6.19	6.22	6.28	6.23	N/A	
emconfiguration	6.19	6.22	6.28	6.23	6.23	
failovermode	6.19	6.22	6.28	6.23	N/A	
faults -list	6.19	6.22	6.28	6.23	6.23	
firmware	6.19	6.22	6.28	N/A	N/A	
flash -ioport	N/A	N/A	6.28	N/A	N/A	
flash -iomodule	N/A	N/A	6.28	N/A	N/A	
flashleds	6.19	6.22	6.28	6.23	6.23	
getagent	6.19	6.22	6.28	6.23	6.23	
getall	6.19	6.22	6.28	6.23	6.23	
getarrayuid	6.19	6.22	6.28	6.23	6.23	
getcache	6.19	6.22	6.28	6.23	6.23	
getconfig	6.19	6.22	6.28	6.23	6.23	
getcontrol	6.19	6.22	6.28	6.23	6.23	
getcrus	6.19	6.22	6.28	6.23	6.23	
getdisk	6.19	6.22	6.28	6.23	6.23	
getlog	6.19	6.22	6.28	6.23	6.23	
getloop	6.19	6.22	6.28	6.23	6.23	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
getlun	6.19	6.22	6.28	6.23	6.23	
getresume	6.19	6.22	6.28	6.23	6.23	
getrg	6.19	6.22	6.28	6.23	6.23	
getsniiffer	6.19	6.22	6.28	N/A	N/A	This command does not support thin LUNs.
getsp	6.19	6.22	6.28	N/A	N/A	
getsptime	6.19	6.22	6.28	N/A	N/A	
getunusedluns	6.19	6.22	6.28	N/A	N/A	
inserttestevent	6.19	6.22	6.28	N/A	N/A	
ioportconfig -list	N/A	N/A	6.28	N/A	N/A	
ioportconfig -persist	N/A	N/A	6.28	N/A	N/A	
luncache -clear	6.19	6.22	6.28	6.23	6.23	This command does not support thin LUNs.
luncache -list	6.19	6.22	6.28	6.23	6.23	This command does not support thin LUNs.
managedby	6.19	6.22	6.28	N/A	N/A	
managefiles -delete	6.19	6.22	6.28	N/A	N/A	
managefiles -list	6.19	6.22	6.28	N/A	N/A	
managefiles -retrieve	6.19	6.22	6.28	N/A	N/A	
ndu -list	6.19	6.22	6.28	N/A	N/A	
ndu -install	6.19	6.22	6.28	N/A	N/A	
ndu -runrules	6.19	6.22	6.28	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
ndu -commit	6.19	6.22	6.28	N/A	N/A	
ndu -status	6.19	6.22	6.28	N/A	N/A	
ndu -revert	6.19	6.22	6.28	6.23	N/A	
networkadmin -get	6.19	6.22	6.28	6.23	6.23	
networkadmin -mib	6.19	6.22	6.28	6.23	6.23	
networkadmin -route	N/A	N/A	6.29	N/A	N/A	
networkadmin -set	6.19	6.22	6.28	6.23	6.23	
port	6.19	6.22	6.28	N/A	N/A	
powersaving	N/A	N/A	6.29	N/A	N/A	
rebootSP	6.19	6.22	6.28	6.23	6.23	
rebootpeerSP	N/A	24	6.28	6.23	6.23	AX150 series also supports this command.
register	6.19	6.22	6.28	N/A	N/A	
remoteconfig -getconfig	6.19	6.22	6.28	N/A	N/A	Not supported on AX4-5 series
remoteconfig -reloadconfig	6.19	6.22	6.28	N/A	N/A	
remoteconfig -setconfig	6.19	6.22	6.28	N/A	N/A	
removerg	6.19	6.22	6.28	N/A	N/A	
resetandhold	6.19	6.22	6.28	N/A	N/A	
responsetest	6.19	6.22	6.28	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
sc_off	6.19	6.22	6.28	N/A	N/A	
setcache	6.19	6.22	6.28	6.23	N/A	
selfeature	6.19	6.22	6.28	N/A	N/A	
setsniffer	6.19	6.22	6.28	N/A	N/A	This command does not support thin LUNs.
setspstime	6.19	6.22	6.28	N/A	N/A	
setsptime	6.19	6.22	6.28	N/A	N/A	
setstats	6.19	6.22	6.28	N/A	N/A	CX series only
shutdown	6.19	6.22	6.28	N/A	N/A	AX4-5 series and AX series only
shutdownsp	6.19	6.22	6.28	N/A	N/A	AX4-5 series and CX4 series only
shutdownpeersp	6.19	6.22	6.28	N/A	N/A	AX4-5 series and CX4 series only
spcollect -info	6.19	6.22	6.28	N/A	N/A	AX4-5 series and CX4 series only
spcollect -set	6.19	6.22	6.28	N/A	N/A	
spcollect	6.19	6.22	6.28	N/A	N/A	
spportspeed	6.19	6.22	6.28	N/A	N/A	
systemtype	6.19	6.22	6.28	6.23	N/A	AX4-5 series running Navisphere Manager and CX series only
storagegroup	6.19	6.22	6.28	N/A	N/A	
trespass	6.19	6.22	6.28	N/A	N/A	
unbind	6.19	6.22	6.28	N/A	N/A	
unitserialnumber	6.19	6.22	6.28	N/A	N/A	
metalun-destroy	6.19	6.22	6.28	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
metalun -expand	6.19	6.22	6.28	N/A	N/A	
metalun -info	6.19	6.22	6.28	N/A	N/A	
metalun -list	6.19	6.22	6.28	N/A	N/A	
metalun -modify	6.19	6.22	6.28	N/A	N/A	
reserved -lunpool -addlun	6.19	6.22	6.28	6.23	6.23	This command does not support the thin LUNs.
reserved -lunpool -list	6.19	6.22	6.28	6.23	6.23	This command does not support the thin LUNs.
reserved -lunpool -rmlun	6.19	6.22	6.28	6.23	6.23	This command does not support the thin LUNs.
connection -adduser	6.19	6.22	6.28	6.23	6.23	
connection -deleteuser	6.19	6.22	6.28	6.23	6.23	
connection -getuser	6.19	6.22	6.28	6.23	6.23	
connection -getport	6.19	6.22	6.28	6.23	6.23	
connection -setport	6.19	6.22	6.28	6.23	6.23	
connection -delport	N/A	N/A	6.29	N/A	N/A	
connection -route	N/A	N/A	6.29	N/A	N/A	
connection -pingnode	6.19	6.22	6.28	6.23	6.23	
connection -traceroute	6.19	6.22	6.28	6.23	6.23	
connection -setsharedauth	6.26	6.22	6.28	6.23	6.23	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
connection -delsharedauth	6.26	6.22	6.28	6.23	6.23	
connection -getsharedauth	6.26	6.22	6.28	6.23	6.23	
connection -addset	6.26	6.22	6.28	6.23	6.23	
connection -delset	6.26	6.22	6.28	6.23	6.23	
connection -modifyset	6.26	6.22	6.28	6.23	6.23	
connection -getset	6.26	6.22	6.28	6.23	6.23	
connection -addpath	6.26	6.22	6.28	6.23	6.23	
connection -delpath	6.26	6.22	6.28	6.23	6.23	
connection -modifypath	6.26	6.22	6.28	6.23	6.23	
connection -verifypath	6.26	6.22	6.28	6.23	6.23	
isns -addserver	6.19	6.22	6.28	6.23	6.23	
isns -deleteserver	6.19	6.22	6.28	6.23	6.23	
isns -listserver	6.19	6.22	6.28	6.23	6.23	
isns -setprimary	6.19	6.22	6.28	6.23	6.23	
migrate -start	6.19	6.22	6.28	6.23	N/A	
migrate -cancel	6.19	6.22	6.28	6.23	N/A	
migrate -modify	6.19	6.22	6.28	6.23	N/A	
migrate -list	6.19	6.22	6.28	6.23	N/A	
domain -add	6.19	6.22	6.28	6.23	N/A	
domain -list	6.19	6.22	6.28	6.23	N/A	



Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
domain -remove	6.19	6.22	6.28	6.23	N/A	
domain -setmaster	6.19	6.22	6.28	6.23	N/A	
ntp -list	6.26	N/A	6.28	N/A	N/A	
ntp -set	6.26	N/A	6.28	N/A	N/A	
security -adduser	6.19	6.22	6.28	6.23	N/A	
security -changeuserinfo	6.19	6.22	6.28	6.23	N/A	
security -list	6.19	6.22	6.28	6.23	N/A	
security -listrole	6.19	6.22	6.28	6.23	N/A	
security -rmuser	6.19	6.22	6.28	6.23	N/A	
security -ldap -addserver	N/A	N/A	6.29	N/A	N/A	
security -ldap -modifyserver	N/A	N/A	6.29	N/A	N/A	
security -ldap -removeserver	N/A	N/A	6.29	N/A	N/A	
security -ldap -listserver	N/A	N/A	6.29	N/A	N/A	
security -ldap -addrolemapping	N/A	N/A	6.29	N/A	N/A	
security -ldap -modifyrolemapping	N/A	N/A	6.29	N/A	N/A	
security -ldap -removeolemapping	N/A	N/A	6.29	N/A	N/A	
security -ldap -listrolemapping	N/A	N/A	6.29	N/A	N/A	
security -ldap -synchronize	N/A	N/A	6.29	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
storagepool -create	N/A	N/A	6.28	N/A	N/A	
storagepool -list	N/A	N/A	6.28	N/A	N/A	
storagepool -destroy	N/A	N/A	6.28	N/A	N/A	
storagepool -modify	N/A	N/A	6.28	N/A	N/A	
storagepool -expand	N/A	N/A	6.28	N/A	N/A	
storagepool -cancelexpand	N/A	N/A	6.28	N/A	N/A	
thinlun -create	N/A	N/A	6.28	N/A	N/A	
thinlun -list	N/A	N/A	6.28	N/A	N/A	
thinlun -modify	N/A	N/A	6.28	N/A	N/A	
thinlun -destroy	N/A	N/A	6.28	N/A	N/A	
thin -info	N/A	N/A	6.28	N/A	N/A	
-portal -create	N/A	N/A	6.29	N/A	N/A	
-portal -list	N/A	N/A	6.29	N/A	N/A	
-portal -migrate	N/A	N/A	6.29	N/A	N/A	
-portal -destroy	N/A	N/A	6.29	N/A	N/A	
-template -list	N/A	N/A	6.29	N/A	N/A	
-template -import	N/A	N/A	6.29	N/A	N/A	
-template -export	N/A	N/A	6.29	N/A	N/A	
-template -destroy	N/A	N/A	6.29	N/A	N/A	
-template -findconflict	N/A	N/A	6.29	N/A	N/A	
-template -swap	N/A	N/A	6.29	N/A	N/A	
-monitor -addsystem	N/A	N/A	6.29	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
-monitor -removesystem	N/A	N/A	6.29	N/A	N/A	
-monitor -listmapping	N/A	N/A	6.29	N/A	N/A	
-monitor -applytemplate	N/A	N/A	6.29	N/A	N/A	
-monitor -stoptemplate	N/A	N/A	6.29	N/A	N/A	
-monitor -getlog	N/A	N/A	6.29	N/A	N/A	
-monitor -enableresponse	N/A	N/A	6.29	N/A	N/A	
-monitor -disableresponse	N/A	N/A	6.29	N/A	N/A	
-monitor -getlogsize	N/A	N/A	6.29	N/A	N/A	
-monitor -setlogsize	N/A	N/A	6.29	N/A	N/A	
-monitor -inserttestevent	N/A	N/A	6.29	N/A	N/A	
-monitor -responsetest	N/A	N/A	6.29	N/A	N/A	
-monitor -reloadconfig	N/A	N/A	6.29	N/A	N/A	
server -volmap	N/A	N/A	6.29	N/A	N/A	
server -update	N/A	N/A	6.29	N/A	N/A	
server -getagent	N/A	N/A	6.29	N/A	N/A	
server -remoteconfig	N/A	N/A	6.29	N/A	N/A	
server -register	N/A	N/A	6.29	N/A	N/A	
hypervisor -addvirtualcenter	N/A	N/A	6.29	N/A	N/A	

Command	CX series	CX3 series	CX4 series	AX4-5 Manager	AX4-5 Express	Comments
hypervisor -modifyvirtualcenter	N/A	N/A	6.29	N/A	N/A	
hypervisor -removevirtualcenter	N/A	N/A	6.29	N/A	N/A	
hypervisor -addESX	N/A	N/A	6.29	N/A	N/A	
hypervisor -modifyESX	N/A	N/A	6.29	N/A	N/A	
hypervisor -removeESX	N/A	N/A	6.29	N/A	N/A	
hypervisor -listESX	N/A	N/A	6.29	N/A	N/A	

## Navisphere Server Utility and Initialization Utility

This appendix provides the CLI commands for Navisphere Server Utility and Navisphere Initialization Utility tools.

- ◆ Using the server utility command line to generate a high-availability report ..... 590
- ◆ Using the command line initialization utility ..... 594

## 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 storage-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 storage 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 Navisphere 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 **cd C:\Program Files\EMC\Navisphere Server Utility** or **cd C:\Program Files (x86)\EMC\Navisphere 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 Navisphere 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 ESX servers, use the following additional switches to specify the ESX server or vCenter for which you want a report on.

**-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 storage system

If you are generating a high-availability report as a prerequisite to update the storage-system software, you need to upload this report to the storage system prior to installing the software on the storage 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 ipaddress -vmuser username -vmpassword password**

where:

**-ip IPAddress**

Specifies the IP address of an SP on the storage 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 storage 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 storage 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 storage 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 storage 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 Navisphere 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 Navisphere 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 storage-system software** - If you are about to perform a storage-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.

## Using the command line initialization utility

For Fibre Channel storage systems, use the Navisphere® initialization utility to discover storage systems and set network parameters (IP address, subnet mask, and default gateway). In addition, for iSCSI storage 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 a Navisphere off-array management station. The only requirement is that the host be connected to the same network subnet as the storage-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 spA | ipaddress | -sphosta spA | hostname] [-ipb spB | ipaddress  
| -sphostb spB | hostname] [-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 storage systems.

### **-all**

Discovers and displays a list of initialized and partially initialized storage systems.

**-xml**

Specifies that the output is in xml format.

**configure**

Specifies the network parameters for the specified storage system.

**-serial** *serialnumber*

Specifies the serial number of the storage system you want to initialize.

**-file** *filename*

Specifies the name of the file that will store all the network parameters.

**-ipa** *spAipaddress* | **-sphosta** *spAhostname*

Specifies the IP address of SP A, or the server name for SP A for the specified storage system.

**-ipb** *spBipaddress* | **-sphostb** *spBhostname*

Specifies the IP address of SP B, or the server name for SP B for the specified storage system.

**-mask** *subnetmask*

Specifies the subnet mask for the specified storage system.

**-gateway** *gateway*

Specifies the gateway for the specified storage system.

**-user** *username*

Specifies the storage-system login username.

**-password** *password*

Specifies the storage-system login password.

**-storagename** *storagename*

Specifies the name for the specified storage 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 below for details.

**A**

- Access Logix option
  - defined 20
- account, user 39
- addpath 404
- Agent
  - configuring 271, 273
- alpa -set 57
- arguments, spaces in 45
- arrayconfig 61, 63
- arrayconfig -duplicate 63
- AX series support 33

**B**

- backendbus -analyze 67
- backendbus -get -all 75
- backendbus -get -physical 74
- backendbus -get -sfpstate 73, 74
- backendbus -get -speeds 71
- backendbus -resetspeed 77
- baseuuid command 79
- basic commands, Chapter 2
  - error codes 540, 578, 590
- battery backup units (BBUs), state 156, 212
- battery test time 156, 291, 299
- bind signature, disk 178
- bind state, disk 174
- bind subcommand (RAID group) 80
- binding restrictions and recommendations 81

**C**

- cache information
  - getting (-getall subcommand) 129
  - setting
    - RAID 3 285
    - read 285
    - write 286
- cachecard subcommand
  - initialize 90
  - list 91
- case (of characters in CLI commands) 39
- chglun subcommand 93
- chgrg subcommand 99
- CLARiiON Open storage-system type 314
- clarstats subcommand 102
- clearing RAID SP error log 101
- clearlog subcommand 101
- CLI
  - Java
    - basic commands 41
    - requirements 39
    - spaces in arguments 45
    - using security 39
- CLI functionality
  - Secure 38
- command line format, spaces in 45
- committing driver packages 237
- Compaq storage-system type 314
- configuration server, defined 31
- configuring Agent 271, 273
- connection -adduser command 374
- connection commands
  - addpath 404

- adduser 374
- deleteuser 376, 396, 399
- delpath 406
- delset 399
- delsharedauth 395
- getport 380
- getuser 378
- modifypath 407
- modifyset 400
- pingnode 389
- setport 384
- traceroute 391
- verifypath 409
- connection -deleteuser command 376, 395, 396, 399
- connection -getport command 380
- connection -getuser command 378
- connection -pingnode command 389
- Connection set database commands
  - addpath 404
  - delpath 409
  - delset 399
  - delsharedauth 395
  - getset 402
  - getsharedauth 396
  - modifyset 400
- connection -setport command 384
- connection -setsharedauth 393, 395
- connection -traceroute command 391
- connection -verifypath 409
- createreg subcommand 105
- CRUs, state 168

## D

- Dell storage-system type 314
- delpath 406
- delset 399
- device map 193
- disk 176
  - errors 176
  - information (-getall subcommand) 133
  - RAID type of (-type) 178
  - vendor ID 179
  - write requests/retries 179
  - write retries 179
- DMP, delay for (ndu -delay) 233

## Domain commands

- add 434
- list 436
- remove 437
- setmaster 439

## driver packages

- committing 237
- installing 230
- listing 228
- reverting 239

## E

- element size, stripe 193
- Empty state, disk 174
- Enabled state, disk 174
- Equalizing state, disk 174
- error
  - codes for basic commands 540, 578, 590
  - codes for MirrorView subcommands 553
  - codes for shared storage commands 551
  - codes for SnapView subcommands 553
  - SP log, clearing 101
- Event Monitor
  - log (getlog) 182
  - response, testing 277

## F

- Failed state, disk 174
- failover types
  - failovermode subcommand 109
  - storagegroup command failovermode argument 313, 314
- failovermode subcommand 109
- fan, state 169
- faults -list 111
- file, security 39
- firmware subcommand 113
- firmware, RAID controller, updating 113
- FLARE software revision (getagent) 120
- flash -iomodule 117
- flash -ioport 115
- flashleds command 117, 118
- Formatting state, disk 174
- Fujitsu storage-system type 314

**G**

getagent subcommand 117, 119  
 getall subcommand 122  
 getarrayuid subcommand 156  
 getcache subcommand 158  
 getconfig subcommand 163  
 getcontrol subcommand 165  
 getcrus subcommand 168  
 getdisk subcommand 174  
 getlog subcommand 182  
 getloop subcommand 184  
 getlun subcommand 186  
 getrg subcommand 203  
 getset 402  
 getsniffer subcommand 207  
 getsp subcommand 210  
 getsptime subcommand 212, 213  
 getuid subcommand 156  
 getunusedluns 213  
 global administrator account 430

**H**

hard errors 176  
 HBA
 

- diagnosing 255
- removing registration 255
- unregistering 255

 hot spare parameters 81  
 hot spare ready state 175  
 HP storage-system type 314  
 hypervisor -addESX 535  
 hypervisor -addvirtualcenter 532  
 hypervisor -listESX 538  
 hypervisor -modifyESX 536  
 hypervisor -modifyvirtualcenter 533  
 hypervisor -removeESX 537  
 hypervisor -removevirtualcenter 534

**I**

individual disk unit parameters 81  
 initialization utility
 

- about 22
- installing 22
- using 22

 inserttestevent subcommand 214

**installing**

driver packages 230  
 pre-installation validation checks 232

**ioportconfig -list 215****ioportconfig -persist 218****iSCSI****connection commands**

adduser 374  
 deleteuser 376, 395, 396, 399  
 getport 380  
 getuser 378  
 pingnode 389  
 setport 384  
 setsharedauth 393  
 traceroute 391

**iSNS commands**

addserver 432, 434, 439, 452

**L**

LAN introduced 367

LCC (link control card) state 170

LEDs, flashing (flashleds command) 117, 118

Licensed Internal Code (LIC), see FLARE

software

lightweight directory access protocol (LDAP) 429

listing installed driver packages 228

loop (Fibre Channel) information (getloop) 184

LUN (logical unit)

binding parameters for RAID types 81

binding restrictions and recommendations  
for RAID types without RAID  
groups 80

caching

prefetch parameters 191  
write 194

changing parameters 93

getting information (getlun) 186

information

-getall subcommand 135

lunmapinfo command 222

private 191

RAID type 194

read cache prefetch parameters 193

read caching, prefetch type parameter 96

statistics (getlun subcommand switches) 186

stripe element size 193

- virtual 306, 307
- write retries 195
- LUN migration 419
- luncache subcommand
  - clear 215, 218, 219
  - list 221
- lunmapinfo subcommand 222
- LUNs (logical units), expanded, *see* metaLUNs

**M**

- managedby subcommand 222
- managefiles (Secure/Classic CLI)
  - delete 223
  - list 225
  - retrieve 226
- memory partitions
  - FC5400/5500 285
  - RAID 3 285
  - read cache 285
  - write cache 286
- metaLUN subcommands
  - destroy 331
  - expand 331
  - info 341
  - list 343
  - modify 350
- metaLUNs
  - defined 328
  - expanding SnapView, MirrorView and SAN Copy LUNs 352
  - overview 328
  - using added capacity
    - for SnapView LUNs 352
    - LUNs in clone groups 352
    - LUNs in mirrors 353
    - LUNs in SAN Copy sessions 353
- migrate subcommands 419
- migrate subcommands, -cancel 423, 523, 526, 527, 528, 531
- migrate subcommands, -list 425
- migrate subcommands, -modify 424
- migrate subcommands, -start 420
- MirrorView
  - error codes 553
  - information (-getall subcommand) 155
- mode, failover 109, 313, 314

- modifypath 407
- modifyset 400
- monitor -addsystem 506
- monitor -applytemplate 509
- monitor -disableresponse 513
- monitor -enableresponse 512
- monitor -getlog 511
- monitor -getlogsize 514
- monitor -inserttestevent 516
- monitor -listmapping 508
- monitor -reloadconfig 520
- monitor -removesystem 507
- monitor -responsetest 517
- monitor -setlogsize 515
- monitor -stoptemplate 510

**N**

- navicli.jar CLI command 41
- Navisphere Initialization Utility, *see* Initialization Utility
- Navisphere security 39
- ndu (non-disruptive upgrade) subcommand
  - commit 237
  - install 230
  - list 228
  - pre-installation validation checks 232
  - reverting 239
  - runrules 235
  - status 238
- networkadmin -get subcommand 241
- networkadmin -set subcommand 246, 248, 251, 386, 387
- node, LAN, defined 367
- Not Connected state, SP 171
- ntp 441, 443
- ntp -list 441
- ntp -list (not supported on AX-series) 441
- ntp -set 443

**O**

- Off state, disk 174
- offset, alignment (bind -offset switch) 86
- operational state
  - CRUs 168
  - See also* state



**P**

password 39  
 performance statistics, logging 292  
 Port ID SCSI ID, setting 57  
 port subcommand 255  
 -portal -create 496  
 -portal -destroy 499  
 -portal -list 497  
 -portal -migrate 498  
 Powering Up state, disk 174  
 powersaving 268  
 prefetch parameters (read cache, getlun) 191  
 prefetch type parameter 96  
 Present state, SP 171  
 private LUN 191  
 PROM revision, SP 120  
 PVLINKS feature (HP-UX) 321

**R**

RAID 3 LUN  
   SP memory partition 285  
 RAID group  
   binding 80  
   changing properties 99  
   creating 105  
   information (-getall subcommand) 147  
   information, getting (getrg) 203  
   removing 275  
 RAID type  
   binding parameters 81  
   disk (-type) 178  
   LUN 194  
 read cache memory 285  
 read error, disk 176  
 read memory partition 285  
 Ready state  
   disk 174  
   hot spare 175  
 rebootpeerSP subcommand 270  
 rebootSP subcommand 269  
 Rebuilding state, disk 174  
 remoteconfig -getconfig subcommand 271  
 remoteconfig -setconfig subcommand 273  
 Removed state  
   disk 175  
   SP 171

removerg subcommand 275  
 requirements, Java CLI 39  
 reserved LUN pool  
   estimating size 357  
 resetandhold 276  
 responsetest subcommand 277  
 reverting driver packages 239  
 revision  
   LCC 170  
   SP 210

**S**

sc\_off subcommand 281  
 SCSI ID of SP port 57  
 Secure CLI about 38  
 Security commands  
   adduser 446  
   changeuserinfo 448  
   list 450  
   rmuser 452  
 security -ldap -addrolemapping 462  
 security -ldap -addServer 453  
 security -ldap -listrolemapping 466  
 security -ldap -listserver 460  
 security -ldap -modifyrolemapping 464  
 security -ldap -modifyserver 456  
 security -ldap -removeolemapping 465  
 security -ldap -removeserver 459  
 security -ldap -synchronize 467  
 security, Navisphere 39  
 segment size/multiplier field 193  
 serial number, SP 210  
 server -getagent 527  
 server -register 531  
 server -remoteconfig 528  
 server -update 526  
 server -volmap 523  
 setcache subcommand 283  
 setsniffer subcommand 288  
 setspstime subcommand 290  
 setsptime subcommand 299  
 setstats subcommand 292  
 SGI storage-system type 314  
 shared storage  
   command error codes 551  
   defined 20

- shutdown 295
- shutdown subcommand 294, 296, 478, 479, 481, 498, 499, 500, 501, 502, 503, 504, 505
- Shutdownpeersp 296
- Shutdownsp 295
- snapshot, adding to storage group (storagegroup command) 307
- SnapView
  - displaying copy LUNs 192
  - information (-getall subcommand) 148
- SnapView error codes 553
- software installation status 238
- SP (storage processor)
  - cache setting 283
  - configuration information 165
  - information 165
  - log (getlog) 182
  - PROM revision 120
  - rebooting 269, 270
  - restarting 269, 270
  - revision 210
  - state 171
  - write cache high watermark 283
  - write cache low watermark 284
- spcollect -set subcommand (automated) 298
- spcollect subcommand (manual) 299
- SPS (standby power supply)
  - setting test time 156, 291, 299
  - state 171
- standby state, disk 175
- state
  - disk 174
  - fan 169
  - LCC (link control card) 170
  - SP 171
  - SPS 171
  - VSC 172
- statistics logging 292
  - resetting 102
- status, ndu subcommand 238
- storage system
  - cache information (-getall subcommand) 129
  - information (-getall subcommand) 125, 148
  - information, general (-getall subcommand) 122
  - overview of configuration and management 31
  - type (systemtype) 321
  - types 314
    - unique ID (UID), getarrayuid command 156
- storagepool -cancelExpand 482
- storagepool -create 473
- storagepool -destroy 478, 498
- storagepool -expand 481, 500, 501, 502, 503, 504
- storagepool -list 475, 497
- storagepool -modify 479, 498
- storage-system caching
  - read cache 285
  - write cache
    - high watermark 283
    - low watermark 284
- stripe element size 193
- subcommands
  - alpha -set 57
  - backendbus -analyze 67
  - backendbus -get -all 75
  - backendbus -get -sfpstate 73, 74
  - backendbus -get -speeds 71
  - backendbus -resetspeed 77
  - baseuuid 79
  - bind (RAID group) 80
  - cachecard -clear 215, 218, 219
  - cachecard -initialize 90
  - cachecard -list 91, 221
  - chglun 93
  - chrg (RAID group) 99
  - clearlog 101
  - clearstats 102
  - createrg (RAID group) 105
  - failovermode 109
  - faults -list 111
  - firmware 113
  - flashleds 117, 118
  - getagent 117, 119
  - getall 122
  - getarrayuid 156
  - getcache 158
  - getconfig 163
  - getcontrol 165
  - getcrus 168
  - getdisk 174
  - getlog 182
  - getloop 184
  - getlun 186

getrg 203  
 getsniffer 207  
 getsp 210  
 getsptime 212, 213  
 insertestevent 214  
 lunmapinfo 222  
 managedby 222  
 managefiles -delete (Secure/Classic CLI) 223  
 managefiles -list (Secure/Classic CLI) 225  
 managefiles -retrieve (Secure/Classic CLI) 226  
 networkadmin -get 241  
 networkadmin -set 246, 248, 251, 386, 387  
 port 255  
 rebootpeerSP 270  
 rebootSP 269  
 remoteconfig -getconfig 271  
 remoteconfig -setconfig 273  
 removevg (RAID group) 275  
 resetandhold 276  
 responsetest 277  
 sc\_off 281  
 setcache 283  
 setsniffer 288  
 setspstime 290  
 setsptime 299  
 setstats 292  
 shutdown 294, 296, 478, 479, 481, 498, 499, 500, 501, 502, 503, 504, 505  
 spcollect (manual) 299  
 spcollect -set (automated) 298  
 storagegroup 305  
 systemtype 321  
 trespass 323  
 unbind 324  
 unitserialnumber 325  
 subcommands, snapshots CLI  
   addluntocache 359  
 Sun Cluster, port 80 behavior mode 325  
 systemtype subcommand 321

## T

-template -destroy 503  
 -template -export 502  
 -template -findconflict 504

-template -import 501  
 -template -list 500  
 -template -swap 505  
 testing an Event Monitor template 214  
 thin -info 491  
 thinlun -create 485  
 thinlun -list 485  
 thinlun -modify 489  
 time  
   setting on SPs 212, 213  
   SP 299  
 trespass  
   subcommand 323  
   types  
     failovermode command 109  
     storagegroup failovermode command 313, 314  
 Tru64 storage-system type 314

## U

unbind subcommand 324  
 Unbound state, disk 175  
 unique ID (UID)  
   example, getlun 196  
   getarrayuid command 156  
   storage system 156  
 unique user ID (Tru64) 79  
 unitserialnumber command 325  
 unshared storage systems, defined 21  
 user, security 39  
 username 39

## V

VDPP80 (port 80) behavior mode 325  
 verifypath 409  
 virtual LUN number 306, 307  
 VMware ESX Server  
   binding LUNs 54  
   metaLUNs 332  
   restrictions 28  
 VSC (power supply) state 172  
 VSS Provider  
   managedby subcommand 222

**W**

World Wide Name as unique ID 46

write cache

    high watermark field 283

    low watermark field 284

    memory partition 286

    state (enabled/disabled) 194

write errors, disk 176

write requests, disk 179

write retries, disk 179