



EMC® MirrorView™ /Asynchronous

Command Line Interface (CLI) Reference

P/N 300-001-335

REV A13

EMC Corporation

Corporate Headquarters:

Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright © 2004 - 2010 EMC Corporation. All rights reserved.

Published July 2010

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

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

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

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

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

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

Corporate Headquarters: Hopkinton, MA 01748-9103

Preface	7
Chapter 1: About EMC MirrorView/A CLI	13
MirrorView/A overview.....	14
Prerequisites.....	14
Configuration guidelines.....	15
Sample configuration.....	16
MirrorView connection requirements.....	17
Cable connections between SPs at the MirrorView sites.....	17
MirrorView features and benefits.....	20
Provision for disaster recovery with minimal overhead.....	21
CLARiiON MirrorView environment.....	21
Bidirectional mirroring.....	21
Integration with EMC SnapView software.....	22
Integration with EMC SAN Copy software.....	22
Replication over long distances.....	23
Application integration.....	23
Using online help.....	23
Chapter 2: MirrorView/A Remote Mirroring Commands	25
MirrorView/A operations overview.....	27
About Secure CLI.....	29
Getting started with Secure CLI.....	30
Using Secure CLI.....	30
naviseccli.....	32
MirrorView/A CLI functions.....	36
mirror -async -addimage.....	38

mirror -sync -allocatelog.....	41
mirror -async -change.....	43
mirror -async -changeimage.....	45
mirror -async -create.....	48
mirror -async -destroy.....	50
mirror -disablepath.....	52
mirror -enablepath.....	54
mirror -async -fractureimage.....	56
mirror -async -info.....	58
mirror -async -list.....	60
mirror -async -promoteimage.....	63
mirror -async -removeimage.....	66
mirror -async-sync -setfeature.....	68
mirror -async -syncimage.....	70
Chapter 3: MirrorView/A Consistency Group Commands.....	73
mirror -async -creategroup.....	74
mirror -async -destroygroup.....	77
mirror -async -addtogroup.....	79
mirror -async -removefromgroup.....	81
mirror -async -change group.....	83
mirror -async -syncgroup.....	85
mirror -async -fracturegroup.....	87
mirror -async -promotegroup.....	89
mirror -async -listgroups.....	91
Appendix A: Responding to Failures.....	95
How MirrorView/A handles failures.....	96
Access to the primary SP fails.....	96
Primary image fails.....	96
Failure of the secondary image.....	98
Promoting a secondary image when there is no failure.....	99
Recovering from serious errors.....	103
How consistency groups handle failures.....	105
Access to the SP fails.....	105
Primary storage system fails.....	105
iSCSI-related failure conditions.....	106

Appendix B: Error Codes	107
Error codes.....	108
Error codes - Hexadecimal.....	113
Appendix C: Secure CLI Command Coverage	115
Secure CLI command coverage.....	116
Glossary	117
Index	121

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

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

About this manual

This manual describes the CLI commands you use to set up, configure, and manage EMC® MirrorView™/Asynchronous software. Each major section includes introductory and format information.

This manual refers to EMC MirrorView/Asynchronous as *MirrorView/A*.

Audience

This manual is intended for those who will use the `naviseccli` command to create and manage remote mirrors on AX4-5 series, CX4 series, CX3 series, and CX series storage systems with the MirrorView/Asynchronous (MirrorView/A) option. Readers of this guide should be familiar with the following topics:

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

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	Revision 11

Storage system removed	Last revision including the storage system
FC series; C series	Revision 11

Related documentation

Related documents include:

For AX4-5 series, CX4 series, CX3 series, and CX series storage systems, refer to the appropriate version of the *EMC Navisphere Command Line Interface (CLI) Reference* for your software revision.

For the most current management and security content for AX4-5 series, CX4 series, CX3 series, and CX series storage systems, refer to the Unisphere™ help, which is available from the Unisphere UI and in the Technical Documentation and Advisories section of the Powerlink website (<http://Powerlink.EMC.com>).

We recommend that you download the latest information before you run the CLI commands.

- ◆ EMC Unisphere Host Agent and CLI Release Notes
- ◆ EMC MirrorView/Synchronous Release Notes
- ◆ The version of this manual that is applicable to your software revision

Special notice conventions

EMC uses the following conventions for special notices:



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

Important: An important note contains information essential to operation of the software.

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

Hint: A note that provides suggested advice to users, often involving follow-on activity for a particular action.

Typographical conventions

EMC uses the following type style conventions in this document.

Type style	Used for
Normal	<ul style="list-style-type: none"> ◆ Running text ◆ Names of resources, attributes, pools, clauses, functions, and utilities
<i>Italic</i>	<ul style="list-style-type: none"> ◆ Titles of publications (citations) ◆ Variables, in running text
<i>Courier italic</i>	Variables, in syntax diagrams and user input (except Celerra)
Courier bold	Command names, options, and keywords
Helvetica bold	<ul style="list-style-type: none"> ◆ User interface elements (what users specifically select, click, or press) ◆ Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)
<i>Courier</i>	URLs, email addresses, pathnames, filenames, directory names, computer names, links, groups service keys, file systems, command names (in running text), user input (such as commands), and notifications (system output, system messages, etc.)
[]	Optional selections
{ }	Required selections
	Alternative selections. The bar means "or"
...	Nonessential information omitted from an example

Where to get help

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

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

Technical support — For technical support, go to Powerlink and choose **Support**. On the Support page, you will see several options, including one for making a service request. Note that to open a service request, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

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

Please send your opinion of this document to:

`techpubcomments@EMC.com`

About EMC MirrorView/A CLI

This chapter introduces the EMC® MirrorView™/Asynchronous software and the EMC Navisphere® storage-system management configurations and architecture.

Note: If you are using the Classic CLI commands, refer to the earlier version of this manual on the EMC Powerlink® website.

This manual refers to the EMC MirrorView/Asynchronous product as MirrorView/A.

Note: If you already familiar with MirrorView/A, you can skip to Chapter 2.

Major topics are:

- ◆ [MirrorView/A overview on page 14](#)
- ◆ [Prerequisites on page 14](#)
- ◆ [Configuration guidelines on page 15](#)
- ◆ [MirrorView connection requirements on page 17](#)
- ◆ [MirrorView features and benefits on page 20](#)
- ◆ [Using online help on page 23](#)

MirrorView/A overview

EMC® MirrorView™/A lets you periodically update a remote copy of production data. It is a software application that keeps a point-in-time copy of a logical unit number (LUN) and periodically replicates the copy to a separate location in order to provide disaster recovery, that is, to let one image continue to be active if a serious accident or natural disaster disables the other. It provides data replication over long distances (hundreds to thousands of miles). To provide for disaster recovery, the primary and secondary storage systems should be geographically separated. MirrorView/A ensures that data from the primary storage system replicates to the secondary.

The production image (the one mirrored) is called the primary image; the copy image is called the secondary image. MirrorView/A supports one remote image per primary. The primary image receives I/O from a server called the production server; a separate storage system maintains the secondary image. This storage system can optionally have a failover/standby computer connected to it or can be connected to its own computer system. Both storage systems can be in different domains, which you manage with the user interface (UI). The client that is managing the storage system containing the primary images can fail over to the secondary image if the primary image becomes inaccessible. After initial synchronization, the remote site always has a consistent point-in-time copy of the primary data.

Important: As a storage-system-based mirroring product, MirrorView/A does not flush server buffers before replicating the primary data. Therefore, the copy is a crash-consistent image of the primary data. You must verify data integrity of the secondary image before using it for disaster recovery. (The verification process varies by application type.)

MirrorView/A supports MirrorView/A consistency groups, which this manual refers to as *consistency groups*. A consistency group is a set of asynchronous mirrors whose secondary images need to be kept consistent with each other in order to be useful; that is, the data on the set of secondary images must have existed on the set of primary images previously. This allows an application to use the secondary images if the primary storage system fails.

Important: The primary images of mirrors in a MirrorView/A consistency group must reside on a single storage system, and the secondary images of the mirrors must reside on a single (but different) storage system. This contrasts with volumes in a Symmetrix® consistency group, which can reside on multiple storage systems.

Prerequisites

- ◆ You must have the MirrorView/A software installed and enabled on all AX4-5 series, CX4 series, CX3 series, and CX series storage systems you want to participate in a mirror. See the Unisphere™ help for information on installing software on the AX4-5 series, CX4 series, CX3 series and CX series storage systems.
- ◆ Data access control must be enabled.

- ◆ You must have Unisphere installed and enabled.
- ◆ SAN configurations must have qualified switches.
- ◆ WAN configurations must have qualified FC-to-IP devices.

Configuration guidelines

The following are configuration rules for MirrorView/A:

- ◆ Each mirror must have one primary image and zero or one secondary image. Any single storage system can have only one image of a mirror.
- ◆ A storage system can have concurrent mirroring connections to a maximum of four other storage systems. (Mirroring connections are common between synchronous and asynchronous mirrors.)
- ◆ The following table lists the configuration limits for the supported platforms.

Platform	CX4 series systems	CX700, CX3 model 40, CX3 model 80 systems	CX500, CX3 model 20 systems	AX4-5 series, CX3 model 10 systems
Primary or secondary mirror images per storage system	256	100	50	25

Note: A metaLUN is a single entity; therefore, it is counted as one of your images. For example, if a mirrored metaLUN is composed of five components, it is counted as one of your images, not five. These limits are independent of the limits for synchronous mirrors. (See the *EMC MirrorView/Synchronous Command Line Interface (CLI) Reference*.)

- ◆ To manage remote mirror configurations, the management workstation must have an IP connection to both the local and remote storage systems. The connection to the remote storage system should have an effective bandwidth of at least 128 Kb/second. The storage systems can be in different domains, which you manage with the UI (see the Unisphere online help).
- ◆ MirrorView/A supports the new, larger LUNs that the EMC FLARE® operating environment supports. (See the FLARE release notes.)
- ◆ Navisphere version 6.29.00 contains thin support for MirrorView/A.
- ◆ You cannot combine thin LUNs from CX4 storage systems running FLARE version 04.29.000.5.xxx or later with pre-FLARE version 04.29.000.5.xxx traditional LUNs or with traditional LUNs from an uncommitted FLARE version 04.29.000.5.xxx. Therefore, you cannot mirror between LUNs created on storage systems running FLARE version 04.29.000.5.xxx and LUNs created on storage systems running pre-FLARE version 04.29.000.5.xxx.

Note: If your system will include Enterprise Flash Drives (solid state disk drives with flash memory, or SSD drives), be sure to carefully plan your configuration using the *Best Practices and planning* documentation available on EMC Powerlink: Home ► Support ► Technical Documentation and Advisories ► White Papers > Configuration/Administration.

Sample configuration

Figure 1 on page 16 shows a sample remote mirror configuration using either iSCSI networks or Fibre Channel switch fabrics. The configuration has two sites and a primary and secondary image that includes the database of four LUNs.

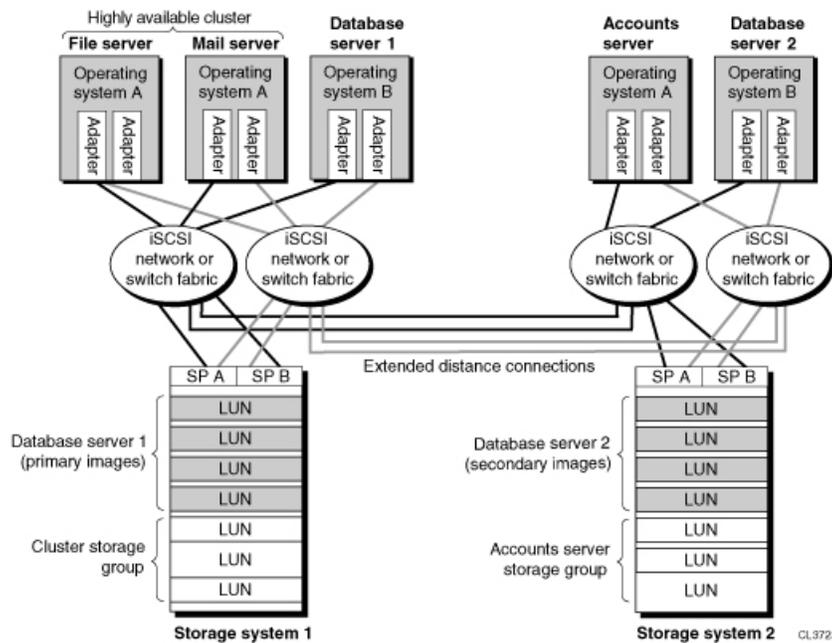


Figure 1. Sample remote mirror configuration

In the above figure, database server 1, which is the production server, executes customer applications. These applications access data on storage system 1, in the database server storage group. Storage system 2 mirrors the data on the database server storage group.

Note: Storage groups are different from consistency groups, which are described in [Chapter 3](#). For more information about storage groups, see the Unisphere online help.

We recommend that you attach a secondary server to storage system 2, so that if a complete site failure occurs where storage system 1 and database server 1 are located, you can completely fail over to the secondary site, and thus minimize the outage window. The server

at the standby site is not required, but because we recommend it, this example includes it in the overall configuration.

Each server has a path to each SP through each network or fabric to each storage system. If a failure occurs in a path, software installed on the server (for example, EMC PowerPath® software) switches to the path through the other SP and continues accessing the data, transparent to the applications on the server.

The production server sends a write request to an SP in storage system 1, which then writes data to the local LUN. The change to the primary LUN is recorded, and at an interval that you define, all changes are copied to the secondary storage system.

If a failure occurs in storage system 1, an administrator can use the client that is managing the storage system to promote the image on storage system 2 to the role of primary image.

Note: The mirrored data is inaccessible until the secondary image is promoted to a primary image.

Then, the appropriate applications can start on any connected server (here, database server 2) with full access to the data. The mirror can be accessible in minutes, although the time needed for applications to recover will vary.

MirrorView connection requirements

MirrorView requires the following:

- ◆ One server, connected to one of the storage systems (a second server, connected to the other storage system, is optional).
- ◆ A Fibre Channel connection (direct or switch) or an iSCSI connection between the two storage systems, shown on the following pages).

Note: The iSCSI connection is not supported on AX4-5 series storage systems with FLARE version 02.23.050.5.5xx.

Cable connections between SPs at the MirrorView sites

MirrorView uses a front-end port on each storage processor (SP) as a communication channel between the storage systems in a remote mirror configuration. This port is called the mirror port in this document.

[Table 1 on page 18](#) shows the mirror ports for the CX4 series storage systems shipped from the factory without optional I/O modules. The MirrorView ports may vary depending on the type and number of I/O modules in the storage system. The *EMC Navisphere Command Line Interface (CLI) Reference* explains how to determine the ports available for MirrorView on storage systems shipped from the factory with optional I/O modules.

Note: For CX4 series systems, the port numbers listed are logical ports; therefore, these CX4 logical port numbers may not be the same as the physical port numbers that Navisphere displays.

Table 1. Mirror ports for the CX4 series storage systems shipped from the factory without optional I/O modules

Storage system	MirrorView Fibre Channel FE ports		MirrorView iSCSI FE ports	
	Logical port ID	Physical slot and port number	Logical port ID	Physical slot and port number
CX4-120, CX4-240	A-1	slot A0 port 3	A-3	slot A1 port 1
	B-1	slot B0 port 3	B-3	slot B1 port 1
CX4-480, CX4-960	A-3	slot A1 port 3	A-5	slot A4 port 1
	B-3	slot B1 port 3	B-5	slot B4 port 1

Table 2 on page 18 shows the mirror ports for storage systems other than the CX4 series systems. All port numbers referenced are front-end ports.

Table 2. Mirror ports for storage systems other than the CX4 series systems

Model	Fibre Channel MirrorView port	iSCSI MirrorView port
CX500	1	N/A
CX700	3	N/A
CX3-10c	3	1 ^a
CX3-20, CX3-20f, CX3-40, CX3-40f	1	N/A
CX3-20c, CX3-40c	5	3 ^a
CX3-80	3	N/A

Note: For information about determining what your MirrorView port is, see the *EMC Navisphere Command Line Interface (CLI) Reference*.

Although server I/O can share the front-end port with MirrorView, for performance reasons, we strongly recommend that server I/O use the front-end ports that MirrorView is not using.



CAUTION Currently, MirrorView and SAN Copy™ software cannot share the same SP port. Before installing the MirrorView enabler, you must deselect any MirrorView ports that a SAN Copy session is using. Otherwise, any SAN Copy sessions using the MirrorView port will fail.

For MirrorView to work correctly, the SP A mirror port at one site must be connected to the SP A mirror port on the other site and the SP B mirror port at one site must be connected to

^a Mirroring over iSCSI is supported for FLARE version 03.26.xxx.5.xxx or later.

the SP B mirror port at the other site. The connections can be either direct or through a switch fabric.

Direct remote mirror connections

A direct mirror configuration consists of one primary storage system and one secondary storage system. The remote mirror connections must be between:

- ◆ SP A mirror ports on the primary and secondary storage systems
- ◆ SP B mirror ports on the primary and secondary storage systems

Figure 2 on page 19 shows a sample direct remote mirror configuration.

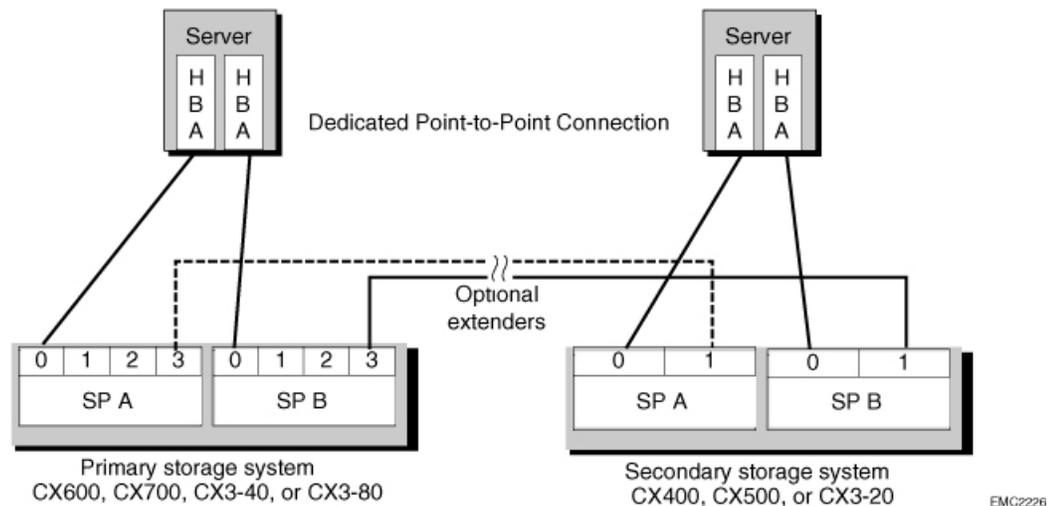


Figure 2. Sample direct remote mirror configuration

Fabric remote mirror connections

A fabric mirror configuration consists of one primary storage system and up to four secondary storage systems. The fabric connections must be as follows:

- ◆ SP A mirror port on the primary storage system must be connected to the same switch fabric as the SP A mirror port on the secondary storage system.
- ◆ SP B mirror port on the primary storage system must be connected to the same switch fabric as the SP B mirror port on the secondary storage system.

Note: The fabric to which SP A mirror ports are connected can be the same fabric or a different fabric than the one to which the SP B mirror ports are connected.

You must zone the mirror port switch connections as follows:

- ◆ A zone for each SP A mirror port on the primary storage system and the SP A mirror port on each secondary storage system.

- ◆ A second zone for each SP B mirror port on the primary storage system and the SP B mirror port on each secondary storage system.

For example, if you have primary storage system 1 and secondary storage systems 2 and 3, you need the following two zones:

- ◆ Zone 1—SP A mirror port on storage system 1 and SP A mirror ports on storage systems 2 and 3.
- ◆ Zone 2—SP B mirror port on storage system 1 and SP B mirror ports on storage systems 2 and 3.

You can use the same SP port for server data and MirrorView. Be careful when an IP distance connection is used because using the same SP port may cause a degradation in both replication and server application performance.



CAUTION MirrorView and SAN Copy software cannot share the same SP port. Before installing the MirrorView enabler, you must deselect any MirrorView ports that a SAN Copy session is using. Otherwise, any SAN Copy sessions using the MirrorView port will fail.

Figure 3 on page 20 shows a sample remote mirror fabric configuration.

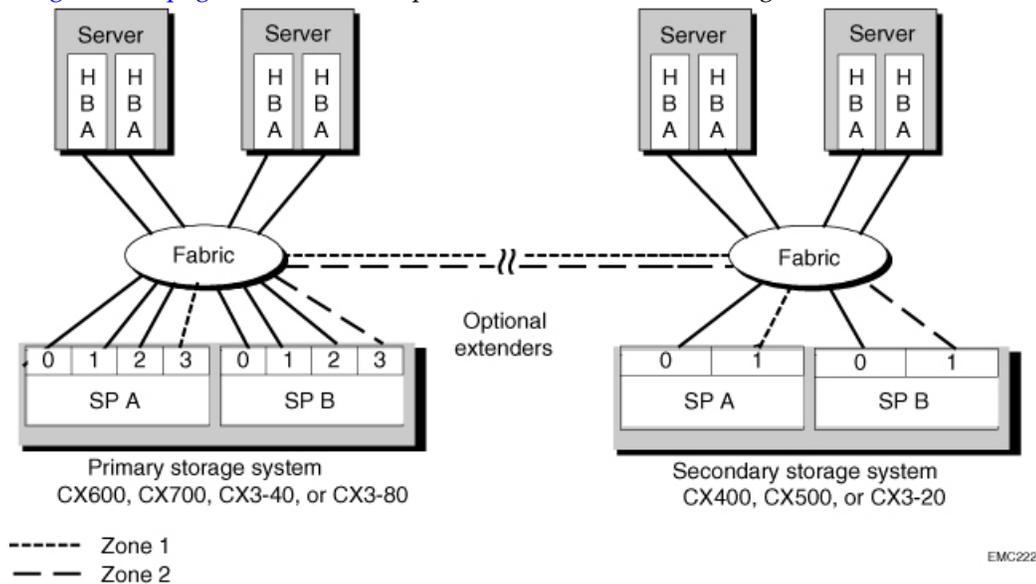


Figure 3. Sample remote mirror connection through a fabric

MirrorView features and benefits

MirrorView mirroring has the following features:

- ◆ Provision for disaster recovery with minimal overhead

- ◆ CLARiiON® environment
- ◆ Bidirectional mirroring
- ◆ Integration with EMC SnapView™ LUN copy software
- ◆ Integration with EMC SAN Copy software
- ◆ Replication over long distances
- ◆ Application integration

Provision for disaster recovery with minimal overhead

Provision for disaster recovery is the major benefit of MirrorView mirroring. Destruction of the data at the primary site would cripple or ruin many organizations. After a disaster, MirrorView lets data processing operations resume with minimal overhead. MirrorView enables a quicker recovery by creating and maintaining a copy of the data on another storage system.

MirrorView is transparent to servers and their applications. Server applications do not know that a LUN is mirrored and the effect on performance is minimal.

With MirrorView/A, secondary systems are periodically updated, based on the user-defined update frequency.

MirrorView is not server-based; therefore, it uses no server I/O or CPU resources. The additional processing for mirroring is performed on the storage system.

CLARiiON MirrorView environment

MirrorView operates in a highly available environment, leveraging the dual-SP design of CLARiiON systems. If one SP fails, MirrorView running on the other SP will control and maintain the mirrored LUNs. If the server is able to fail over I/O to the remaining SP, then periodic updates will continue. The high-availability features of RAID protect against disk failure, and mirrors are resilient to an SP failure in the primary or secondary storage system.

Bidirectional mirroring

A single storage system may be primary (that is, hold the primary image) for some mirrors and secondary (that is, hold the secondary image) for others. This enables bidirectional mirroring.

Note: A storage system can never hold more than one image of a single mirror; that is, it cannot contain both the primary image and secondary image of a single mirror.

Integration with EMC SnapView software

SnapView software lets you create a snapshot of an active LUN at any point in time. The snapshot is a consistent image that can serve for other application purposes while I/O continues to the source LUN. The secondary image is not viewable to any servers, but you can use SnapView in conjunction with MirrorView/A to create a snapshot of a secondary image on a secondary storage system to perform data verification and run parallel processes.

Note: Before taking a snapshot, make sure that the secondary LUN is in a normal condition or fractured, and is not updating. A snapshot of a secondary LUN when MirrorView/A is updating the secondary storage system will not give consistent data. Also, note that data cached on the server, as well as data written to the primary storage system but waiting to be transferred to the secondary storage system on the next update, will not be included in a snapshot of the secondary image.

You can clone either a primary or secondary image by creating a clone group on the same LUN as the mirror image and then adding clones to the group. For more information about using SnapView with MirrorView, see the Unisphere online help.

Note: Before fracturing a clone or starting a SnapView session, make sure that the secondary image is in the synchronized or consistent state. Fracturing a clone or starting a SnapView session of a secondary LUN when MirrorView/A is synchronizing the secondary storage system will not give consistent data. Also, note that data cached on the server, as well as data written to the primary storage system but waiting to be transferred to the secondary storage system on the next update, will not be included in the replica of the secondary image.

The server (if any) connected to the secondary storage system might sit idle until the primary site fails. With SnapView at the secondary site, the server at the secondary site can take snapshots of the mirror images and back them up to other media. This provides point-in-time snapshots of production data with little impact to production server performance.

Integration with EMC SAN Copy software

SAN Copy software lets you create an intra- or inter-storage system copy of a LUN at any point in time. However, do this only when the image state is either Synchronized or Consistent and the mirror is not updating to the secondary image. The copy is a consistent image that can serve for other application purposes while I/O continues to the source LUN. The MirrorView secondary image is not viewable to any servers, but you can use SAN Copy to create a copy of the secondary image on a secondary storage system to perform data verification and run parallel processes.

Note: Related to the process of making the MirrorView secondary unavailable for server I/O, you cannot run SAN Copy full copy sessions on MirrorView secondary images. You can, however, run SAN Copy incremental sessions on MirrorView secondary images.

For more information about SAN Copy, see the Unisphere online help.

Replication over long distances

MirrorView/A uses FC-to-IP devices to provide replication over long distances (hundreds to thousands of miles).

Application integration

Using the Navisphere commands, you can script application integration. You can do this by setting the mirror to the manual update type and issuing a syncimage command from the CLI. To ensure that the application data is in a consistent state, put the application (for example, a database) in backup mode, make sure any data buffered on the server is flushed to the storage system, and issue the MirrorView/A **syncimage** command. After the update has started, you can resume the application activities. These application integration processes allow MirrorView/A to generate a secondary that has a consistent state image of the primary and, thus you can readily use it for backup or application testing.

Using online help

The following online help is available from the Unisphere interface:

- ◆ A set of organized, linked help topics
 - To access the online help table of contents, select **Help ► Help Topics** on the menu bar in the application's main window, or click the help icon in the toolbar.
- ◆ Context-sensitive help topics

To display context-sensitive help, click the **Help** button displayed in each dialog box.

MirrorView/A Remote Mirroring Commands

This chapter explains the `naviseccli` commands and the commands for creating and managing remote mirrors on AX4-5 series, CX4 series, CX3 series, and CX series storage systems with the MirrorView/A option. These commands let you use MirrorView/A software to create a byte-for-byte copy of one or more local LUNs connected to a distant storage-system server.

Note: The commands in this chapter function only with a storage system that has the optional MirrorView/A software installed.

Note: Starting with the 02.26 and 03.26 version of FLARE, Classic CLI can be disabled through the Unisphere UI or a Secure CLI command. Only Secure CLI supports new features. Secure CLI is more secure than Classic CLI and is the preferred interface. Other than commands issued to host agents, the 02.26 and 03.26 versions of Secure CLI includes all the same commands as Classic CLI. If you plan to use Classic CLI commands, you must make sure that Classic CLI is enabled using the Unisphere UI or Secure CLI command. For a list of commands Secure CLI does not support, see the *EMC Navisphere Command Line Interface (CLI) Reference*. If you are currently using Classic or Java CLI to issue CLI commands, note that Secure CLI is replacing both Classic and Java CLI. See *EMC MirrorView/Asynchronous Command Line Interface (CLI) Reference* for Java CLI support.

Major topics are:

- ◆ [MirrorView/A operations overview on page 27](#)
- ◆ [About Secure CLI on page 29](#)
- ◆ [Getting started with Secure CLI on page 30](#)
- ◆ [naviseccli on page 32](#)
- ◆ [MirrorView/A CLI functions on page 36](#)
- ◆ [mirror -async -addimage on page 38](#)
- ◆ [mirror -sync -allocatelog on page 41](#)
- ◆ [mirror -async -change on page 43](#)
- ◆ [mirror -async -changeimage on page 45](#)

- ◆ `mirror -async -create` on page 48
- ◆ `mirror -async -destroy` on page 50
- ◆ `mirror -disablepath` on page 52
- ◆ `mirror -enablepath` on page 54
- ◆ `mirror -async -fractureimage` on page 56
- ◆ `mirror -async -info` on page 58
- ◆ `mirror -async -list` on page 60
- ◆ `mirror -async -promoteimage` on page 63
- ◆ `mirror -async -removeimage` on page 66
- ◆ `mirror -async-sync -setfeature` on page 68
- ◆ `mirror -async -syncimage` on page 70

MirrorView/A operations overview

This section provides an overview of MirrorView/A operations and how to manage the storage systems using the management software.

1. Connect the same management software to both storage systems and configure the software, so that you can manage both storage systems.

Note: Optionally, you can perform this step at the secondary site as well.

You must manage both storage systems, which can be in different domains (see the Unisphere online help).

2. Estimate the size of the reserved LUN pool and the reserved LUNs (see the *EMC Navisphere Command Line Interface (CLI) Reference*).
3. Establish a usable, two-way connection between the MirrorView/A storage systems using the `mirror -enablepath` command (see [mirror -enablepath on page 54](#)).
4. If the primary LUN does not exist, bind it on its server's storage system. Wait for the LUN to finish binding and add it to the storage group.
5. If the secondary LUN does not exist, create a secondary image LUN.

Note: The secondary LUN can be a different RAID type from the primary, as long as the block size matches.

6. Wait for the secondary LUN to finish binding.
7. From the management workstation, create the remote mirror (see [mirror -async -create on page 48](#)).

Next add a secondary image. To add a secondary image, see [mirror -async -addimage on page 38](#).

Normally, when you add a secondary image to a mirror, the software synchronizes the secondary image with the primary image .

The software will initially copy all the data from the primary LUN to the secondary LUN, in order to ensure that the two LUNs are identical.

At any time in the previous sequence of steps, you can get remote mirror status with the `mirror -async -list` command.

8. Create a consistency group and add the primary image to it (optional). See [mirror -async -creategroup](#) on page 74.
9. If a primary failure occurs, Navisphere reports the failure.

If the primary failure is minor, have the primary fixed and resume mirroring.

If the primary failure is catastrophic, the original client that is managing the storage system may be unusable and thus unable to report the failure. For such a failure, the administrator at the secondary site must set up a client to manage the storage system (if not already done), then promote the secondary to primary and take other recovery action as needed. This includes assigning the newly promoted LUN to a storage group, if it is not already the appropriate one.

When you use MirrorView/A on a VMware ESX Server, after you promote the secondary image to a primary image, perform the following steps:

- a. If not already assigned, assign the newly promoted primary image to a storage group on the same or standby ESX Server.
- b. Rescan the bus at the ESX Server level.
- c. If not already created, create a virtual machine (VM) on the same or standby ESX Server. The VM is not powered up when you create it.
- d. Assign the newly promoted primary to the VM.
- e. Power up the VM.

If the VM is created and running and you have not already assigned the newly promoted primary to the VM, perform these steps:

- a. Perform steps a and b above.
- b. If you are running ESX Server 2.5.x, power it down.
- c. To assign the newly promoted primary to the VM, use the virtual center interface for ESX Server 3.x and 2.5.x or the Management User interface for ESX Server 2.5.x.

Note: If you are running ESX Server 3.x, you do not need to power down the VM.

- d. If you are running ESX Server 2.5.x, power up the VM.

Note: If you are running ESX Server 3.x, you do not need to power up the VM.

The primary image (which is now the secondary image) will not be accessible to the primary ESX Server.

Note: For configuration restrictions when using VMFS volumes, go to the E-Lab™ Navigator on the EMC Powerlink website, and under the PDFs and Guides tab, open the `VMware ESX server .pdf` file.

10. If access to the secondary storage image fails, the primary storage system will fracture the remote mirror. If the problem with the secondary is minor (for example, replacing a cable), then the administrator can fix it. Mirroring will recover and resynchronize the image, if the image recovery policy is Automatic.

Whenever you want to stop mirroring, you can first fracture and remove the secondary images, and then destroy the mirror. This does not affect any data on either image LUN, and access to the primary LUN is also unaffected. The LUN that held the secondary image now becomes accessible as a regular LUN.

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. You do not need to install a JRE to run Secure CLI.

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

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

Note: For commands that originated in Classic CLI, some command output may be enhanced; for example, Secure CLI can retrieve and display information from peer SPs. For Classic CLI commands that produce multiple warnings and require multiple confirmations, Secure CLI provides a single summary of warnings and a single confirmation. In general, Secure CLI preserves original command syntax and output for script compatibility.

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 a number 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 user account on the storage system. To create the required user accounts using Navisphere CLI, refer to the *EMC Navisphere Command Line Interface (CLI) Reference*. For details on using Secure CLI, refer to the Unisphere online help.

You can also choose to configure Secure CLI to issue Secure CLI commands on the host (see [Using Secure CLI on page 30](#)). 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: Establishing a security file to use commands that originated in Classic CLI ensures that other than the update from **navicli** to the **naviseccli** command, you do not need to modify any established scripts you may have. Secure CLI, unlike Classic CLI, requires the options **-user**, **-scope**, and **-password** (or the password prompt) in each command line; you do not need to provide these options in the command line if you establish a security file.

Using Secure CLI

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 Unisphere security file.

The Unisphere 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.

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).

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 through the command line.

PREREQUISITES

Anyone that can log in to the server running Navisphere CLI 6.X or later.

DESCRIPTION

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

SYNTAX

```
naviseccli -help
naviseccli [-address IPAddress | NetworkName | -h IPAddress | NetworkName]
[-AddUserSecurity]
[-f filename]
[-m]
[-nopoll | -np]
[-parse | -p]
[-password password]
[-port port]
[-q]
[-RemoveUserSecurity]
[security -certificate]
[-timeout | -t timeout]
[-user username]
[-v]
[-xml]
CMD[optional-command-switches]
```

OPTIONS

-help

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

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

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

-AddUserSecurity

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

-f *filename*

Stores the data in a file.

-m

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

Note: This switch is supported only 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 is automatically issues a poll request unless this switch is specified.

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

-parse | **-p**

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

-password *password*

Specifies the password on the 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 will need to specify the **-port** switch and number 2163 in every subsequent command you issue.

-q

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

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 server.

-scope 0 | 1 | 2

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

-secfilepath *filename*

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

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

security -certificate

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

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

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

-user *username*

Specifies the username on the 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: This switch is supported only for commands that originated in Classic CLI.

-xml

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

CMD

One of a set of commands used with the **naviseccli** command. The CMD switches are described on the pages that follow.

MirrorView/A CLI functions

The Navisphere CLI provides one of the two interfaces to the MirrorView/A mirroring software. The Unisphere UI provides the other.

You can use the `naviseccli mirror` CLI command to set up and manage remote mirror sessions.

The `naviseccli mirror` command functions are summarized in [Table 3 on page 36](#).

Table 3. naviseccli MirrorView/A command functions

Essential functions (in order performed)	
<code>mirror -enablepath</code>	Opens a path between the local and remote storage system. Required regardless of the type of connection between storage systems.
<code>mirror -async -create</code>	Creates a new remote mirror using an existing LUN. The LUN on which the <code>-create</code> command is executed becomes the primary image.
<code>mirror -async -addimage</code>	Adds a secondary image to a mirror using a LUN on a remote storage system. The default is to synchronize unless you specify otherwise.
<code>mirror -async -changeimage</code>	Changes secondary image properties such as auto recovery setting and recovery policy.
<code>mirror -async -fractureimage</code>	Fractures a secondary image, suspending mirroring of data to that image. A fracture log is maintained, so that the image can be quickly resynchronized later.
<code>mirror -async -removeimage</code>	Removes a secondary image from a mirror.
<code>mirror -async -promoteimage</code>	Promotes a secondary image to primary, while simultaneously demoting the primary image (if still accessible) to secondary. Required in order to access the data in the secondary image in the event of a failure of the primary.

Table 3. navisecli MirrorView/A command functions (continued)

Essential functions (in order performed)	
<code>mirror -async -syncimage</code>	Starts synchronization of the secondary image with the primary. Needed if automatic recovery is not on. I/O is allowed with the primary while synchronization occurs. You can specify a synchronization rate to avoid serious performance impact.
<code>mirror -async -info</code>	Displays MirrorView/A information about a storage system.
<code>mirror -async -list</code>	Displays information on existing mirrors.
Optional reconfiguration functions (alphabetically)	
<code>mirror -async -change</code>	Changes remote mirror properties.
<code>mirror -async -destroy</code>	Terminates remote mirroring of a LUN. This does not affect the data in the LUN, and the server continues to access it.
<code>mirror -disablepath</code>	Severs all paths between the local storage-system SPs and the remote storage-system SPs.

mirror -async -addimage

Adds an image to a previously created mirror.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to add a secondary image to a mirror using a LUN on a remote storage system. The `-addimage` function is required if you want mirroring to occur. You must direct the command to the storage system that holds the primary image. This command lets you create primary and secondary images on thin LUNs, if all the participating storage systems support mirroring on thin LUNs.

Note: The Virtual Provisioning™ enabler must be installed to create thin LUNs.

You can use the `-info` function to display the UIDs of storage systems that have images that can be mirrored.

Note: To add a secondary image, you must have a working connection to the remote storage system, and the remote LUN must be exactly the same size (number of blocks) as the primary image.

Note: In this syntax "or" means that you can specify the switch, `-arrayhost sp-hostname | sp-IP-address -lun lun-number | -lunuid lun-uid` or you can specify the switch, `-arrayuid storage-system-uid -lunuid lun-uid`. The rest of the syntax is common to both.

SYNTAX

```
mirror -async -addimage -name name | -mirroruid mirroruid
-arrayhost sp-hostname|sp-IP-address
-lun lun-number | -lunuid lun-uid
```

```
mirror -sync -addimage -name name| -mirroruid mirroruid
-arrayhost sp-hostname -lun lun-number -lunuid lun-uid
[-recoverypolicy auto|manual]
[-syncrate high|medium|low] [-nosyncrequired ]
```

```
mirror -async -addimage -name name | -mirroruid mirroruid
-arrayuid storage-system-uid -lunuid lun-uid
[-recoverypolicy auto|manual]
[-syncrate high|medium|low]
[-manualupdate | -delaystart time | -enddelay time]
```

`[-nosyncrequired]`

OPTIONS

`-name name`

Specifies the name that you gave to the mirror with the `-create` command.

`-mirroruid uid`

Specifies the unique 16-byte primary image ID (World Wide Name) the software assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the [mirror -async -list on page 60](#) function).

`-arrayhost sp-hostname | sp-IP-address`

Specifies the hostname or network address of the remote SP that will hold the secondary image. The secondary image will be added to the remote mirror connected to this host's SP. Use this switch or `-arrayuid`.

`-lun lun-number`

Specifies the LUN number that will serve as the secondary image of the mirror.

`-arrayuid storage-system-uid`

Specifies the unique 8-byte storage system ID of the remote SP. You can determine this by using the `mirror -info` command or the `getarrayuid` command (see the *EMC Navisphere Command Line Interface (CLI) Reference*). For any letters in the ID, use the same case that the software assigned when the mirror was created. Use this switch or `-arrayhost`.

`-lunuid lun-uid`

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to be added as a secondary image. If you specified `-arrayuid`, then `-lunuid` is required. If you specified `-arrayhost`, then you can use either `-lun` or `-lunuid`.

Note: When you specify `-arrayuid` and `-lunuid`, you must add the mirror driver on the remote LUN using the `-setfeature` command before executing the `-addimage` command.

`-recoverypolicy manual | auto`

Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:

`manual`: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

auto: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.

-syncrate high|medium|low

Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect storage-system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other storage system operations.

-manualupdate

Indicates that you must explicitly update the image.

-delaystart *time*

Specifies the time (in minutes) from the beginning of the previous update to the start of the next update. The current update must complete before the next one can start. If the update is still in process when the time period expires, the next update will start immediately once the current update completes.

-enddelay *time*

Specifies the time (in minutes) from the end of the previous update to the next update. A value of 0 causes updates to occur as fast as possible.

-nosyncrequired

Specifies that the image will be added in a synchronized state; thus, no initial synchronization will occur.

EXAMPLE

Any of the following commands adds an image to a mirror. Note that if you specify the remote storage system using the **-arrayuid** switch, before issuing the command you must explicitly add drivers to the LUN stack using the **-setfeature** command. For SP `ss1_spa`, any of the above commands adds to the remote mirror `db_mirror` the LUN in storage system `50:06:06:10:00:FD:A1:69` (or a LUN on the remote storage system `cpc426`) whose unique LUN ID (world wide name, WWN) is `60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11` or whose LUN number is 9. See also the example in [mirror -async -create on page 48](#).

```
naviseccli -h ss1_spa mirror -async -addimage
-name db_mirror -arrayuid 50:06:06:10:00:FD:A1:69
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

```
naviseccli -h ss1_spa mirror -async -addimage
-name db_mirror -arrayhost cpc426 -lun 9
```

```
naviseccli -h ss1_spa mirror -async -addimage
-name db_mirror -arrayhost cpc426
-lunuid 60:06:06:5A:62:02:00:00:92:95:DE:7C:E4:71:D4:11
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -sync -allocatelog

Allocates or re-allocates a LUN for use as a write intent log.

PREREQUISITES

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

DESCRIPTION

The `naviseccli mirror` command with `-allocatelog` assigns a LUN as the write intent log on each SP of the storage system. The recommended (and minimum) size for the log is 128 MB. The write intent log allows recovery with a partial resynchronization if a recoverable failure of the primary SP occurs. If you do not use the intent log, a full synchronization is always required in the event of a failure of the SP controlling the primary image.

You must specify both SPs in one command line.

SYNTAX

```
mirror -sync -allocatelog -spA LUN-number -spB LUN-number [-unbind] [-o]
```

OPTIONS

`-spA LUN-number` and `-spB LUN-number`

Identify the LUN to be used for the write intent log. It must hold at least 128 MB, must not be part of a storage group and must not be a hot spare.

`-unbind`

If write intent log LUNs are currently defined, unbinds the old LUNs after the write intent log function starts using the new LUNs.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command allocates as the write intent log for SP A the LUN with ID `122` and for SP B the LUN with ID `124`.

```
naviseccli -h ss1_spa mirror -sync -allocatelog -spA 122 -spB 124
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stderr`. You may also receive other error messages.

mirror -async -change

Changes mirror properties.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to change one or more parameters of a previously created mirror.

You must direct the command to the storage system that holds the primary image. You must specify at least one of the optional switches with this command. If you omit the override option, then the CLI displays a message for each parameter to confirm the change.

SYNTAX

```
mirror -async -change -name name | -mirroruid uid
[-description description]
[-requiredimages num_of_images] [-o]
[-enablereadonlysecondary yes|no]
```

OPTIONS

-name *name*

Specifies the name the `-create` command gave the mirror. You can specify either the mirror name or unique ID (next).

-mirroruid *uid*

Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

-description *description*

Provides a meaningful title for the mirror (for example, Employee and Benefits Table for Human Resources). The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.

-requiredimages *num_of_images*

Specifies the minimum number of secondary images that should be accessible to the mirror. The default value is 0. Allowed values are 0, 1, and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.

-usewriteintentlog yes | no

Directs the software to use or not use the remote mirror write intent log. If you want to use the log, you must have allocated it with **-allocatelog**.

yes

Use or start using, the write intent log.

no

Do not use, or stop using, the write intent log.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For the remote mirror **dbmirror** on SP **ss0_spa**, this command reduces the minimum number of accessible secondary images to 0. Use the **-list** command to confirm that the changes were made.

```
naviseccli -h ss0_spa mirror -async -change -name
dbmirror -requiredimages 0
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a Command is not supported error message is printed to stdout. You may also receive other error messages.

mirror -async -changeimage

Changes secondary image properties.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to change one or more parameters of a mirror image.

You must direct the command to the storage system that holds the primary image. You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

Do you want to change the sync rate from Medium to High (y/n)?

To change a parameter, enter `y`.

SYNTAX

```
mirror -async -changeimage -name name | -mirroruid uid
-imageuid image-uid
[-recoverypolicy manual|auto]
[[-manualupdate] | [-delaystart time]]
[-enddelay time]
[-syncrate high|medium|low] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the remote mirror in the `-create` command. You can specify either the mirror name or uid (next).

`-mirroruid uid`

Specifies the unique 16-byte ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

`-imageuid image-uid`

Specifies the unique 8-byte ID of the secondary image you want to change. For any letters in the ID, you must specify the same case that the software assigned to the image when the mirror was created (use the `mirror -list` command).

-recoverypolicy manual|auto

Specifies the policy for recovering the secondary mirror image after a system fracture. Values are:

manual: The administrator must explicitly start a synchronization operation to recover the secondary mirror image.

auto: Recovery automatically resumes as soon as the primary image determines that the secondary mirror image is once again accessible. This is the default.

-manualupdate

Indicates that you must explicitly update the image.

-delaystart *time*

Specifies the time (in minutes) from the beginning of the previous update to the start of the next update. The current update must complete before the next one can start. If the update is still in process when the time period expires, the next update will start immediately once the current update completes.

-enddelay *time*

Specifies the time (in minutes) from the end of the previous update to the next update. A value of 0 causes updates to occur as fast as possible.

-syncrate high|medium|low

Specifies a relative value (low, medium, or high) for the priority of completing updates. High completes updates faster, but may significantly affect storage system performance for server I/O requests, particularly where there are multiple concurrent synchronizations occurring. Low completes updates much slower, but also minimizes impact on other storage-system operations.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For storage system `server1`, mirror `db_mirror`, image `50:06:06:10:00:FD:A1:6`, this command changes the synchronization rate to high. Use the `-list` command to confirm that the changes were made.

```
naviseccli -h erver1 mirror -async -changeimage -name
db_mirror -imageuid 50:06:06:10:00:FD:A1:6 -syncrate high
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -create

Creates a new asynchronous mirror using an available LUN for mirroring.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to create a new mirror using an existing LUN. The LUN on which you run the command becomes the primary LUN. But, for mirroring to occur, you must add an image with the `-addimage` function. Later, you can change some of the parameters you specify here with the `-change` function. This command supports thin LUNs after you install the latest Navisphere version, which contains thin LUN support for MirrorView.

If you specify a minimum number of required images (other than the default value zero) in the `-create` command, the mirror will start in the attention state, since you have not yet added the required secondary images. To avoid this problem, do not specify a minimum number of images with `-create`; instead, create the mirror, add the image, and then if you want a minimum number of required images, change the mirror with the `-change` command.

SYNTAX

```
mirror -async -create -name name
-lun lun_number
[-description description]
[-requiredimages num_of_images]
[-o]
```

OPTIONS

`-name` *name*

Lets you assign a meaningful name to the mirror; for example "Employee Benefits". To include spaces in the name, enclose the entire string in quotation marks as in the example. You can use this name to manage the mirror. The name can include as many as 64 characters.

`-lun` *lun_number*

Specifies the LUN number that will serve as the primary image of the mirror you are creating.

`-description` *description*

Lets you assign a description for the mirror. The software does not use the description to identify the mirror. You can specify up to 256 characters, but we suggest that you should not exceed the line length defined by the operating system.

-requiredimages *num_of_images*

Specifies the minimum number of secondary images that should be accessible to this mirror. The default value is 0. Allowed values are 0, 1 and all. If the number of accessible images goes below the specified value, the mirror goes into the Attention state and generates a log message. You can configure this to generate an appropriate warning (for example, e-mail message or page) to the administrator to enable the administrator to correct the problem.

-o

Executes the command without prompting for confirmation.

EXAMPLE

The mirror **-async -create** command creates the mirror named **db_mirror** on LUN 112. The mirror automatically activates.

```
naviseccli -h ssl_spa mirror -async -create -name db_mirror -lun 112
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -destroy

Destroys a mirror.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to destroy a mirror. It is the counterpart of `-create`. It ends mirroring of the LUN that is part of the mirror. This does not affect the data in the LUN, and the server continues to access it.

Note: You cannot destroy a mirror while it is a member of a consistency group.

You must remove the secondary image in the mirror before issuing the command to the storage system holding the primary image. If you cannot remove the secondary image or access the primary image, you can use the `-force` option. You must direct the command to the storage system that holds the primary image unless `-force` is specified, in which case, you can direct it to either storage system.



CAUTION Use `-force` only if all other options to remove an image and the mirror have failed.

If you omit the override option, a message to confirm the destruction of the mirror appears:

```
Do you really want to destroy the mirror <mirror-name> (y/n)?
```

To destroy the mirror, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -async -destroy -name name | -mirroruid uid
[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the mirror with the `-create` command. You can also use the remote mirror ID.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID assigned when the mirror was created. For any letters in the ID, you must specify the same case as the software assigned when the mirror was created (use the `mirror -list` command).

-force

Forcibly destroys the mirror. This option affects only the local storage system.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, this command destroys the remote mirror that was created with the name `db_mirror`.

```
naviseccli -h ss1_spa mirror -async -destroy db_mirror
```

```
Do you really want to destroy the mirror  
db_mirror?(y/n)? y
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages if the command fails for some reason.

mirror -disablepath

Disables the communication paths created by `-enablepath` to remote mirror SPs.

PREREQUISITES

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

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 32](#)) to sever the path established by `-enablepath` in the remote storage system.

When issued to any SP, the command severs paths for both SPs in both the primary and secondary storage systems; that is, it severs these paths if:

- ◆ SP A of the local storage system is connected to SP A of the remote storage system
- ◆ SP A of the remote storage system is connected to SP A of the local storage system
- ◆ SP B of the local storage system is connected to SP B of the remote storage system
- ◆ SP B of the remote storage system is connected to SP B of the local storage system

This means you can issue the command to either SP of either storage system to sever the paths between the storage systems you are targeting.

If you issue the command while a mirror exists, the CLI issues a warning message; if you continue, any mirror connected over the specified path will be fractured.

Note: If both asynchronous and synchronous mirrors are installed, you need to run the command only once. It will take effect for both mirrors.

If you omit the override option, the CLI displays a message to confirm the disabling of the path. To disable the path, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -disablepath connected-SP-hostname [-connectiotype fibre|iscsi]
[-o]
```

OPTIONS

`-connectiotype fibre | iscsi`

Specifies the connection type as fibre or iSCSI for the path you are disabling. If you do not specify the `-connectiotype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

-o

Executes the command without prompting for confirmation. Without this switch, the CLI displays a confirmation message for each SP.

EXAMPLE # 1

For SP **ss1_spa**, this command and confirmation remove the paths between the local SPs and the remote mirror SPs.

```
naviseccli -h ss1_spa mirror -disablepath ss27_spa -connectiontype fibre
```

Disabling of the path will stop mirroring of data to all images if any among these arrays. A system fracture of all mirrors between two arrays on following storage processors will occur due to this action.

```
128.221.39.29
```

```
10.14.20.154
```

```
Do you want to continue (y/n)? y
```

EXAMPLE # 2

The **mirror -info** command displays the updated status.

```
naviseccli -h ss1_spa mirror -async -info
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a Command is not supported error message is printed to stdout.

mirror -enablepath

Establishes a path for remote mirroring between a local and remote storage system.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to create a path between the specified SP and its peer SP to a remote SP and its peer. These paths will serve as paths for remote mirroring. A remote mirror requires a path between the primary and secondary storage systems.

When issued to any SP, the command creates paths for both SPs in both storage systems; that is, it creates these paths:

- ◆ SP A of the local storage system to SP A of the other storage system
- ◆ SP A of the other storage system to SP A of the local storage system
- ◆ SP B of the local storage system to SP B of the other storage system
- ◆ SP B of the other storage system to SP B of the local storage system

The SP As or the SP Bs in both storage systems must be physically connected.

If you cannot enable paths on one SP due to some error, an informational message appears and the paths are enabled for the other SP. To obtain status information, run the `mirror -info -systems` command.

Note: If both asynchronous and synchronous mirrors are installed, you will need to run the command only once. It will take effect for both mirrors.

SYNTAX

```
mirror -enablepath SP-hostname [-connectiotype fibre|iscsi]
```

OPTIONS

`-connectiotype fibre | iscsi`

Specifies the connection type as fibre or iSCSI for the path you are enabling. If you do not specify the `-connectiotype` switch, the system first tries to find a fibre connection; if it does not find one, it looks for an iSCSI connection.

Note: To change the connection type, you must first disable the current enabled path. To modify, test, or delete connections or change the credentials, see the *EMC Navisphere Command Line Interface (CLI) Reference*.

EXAMPLE

For SP **ss1_spa**, this command enables the paths between **ss1_spa** and its peer SP and **ss27_spa** and its peer SP. The **mirror -info** command that follows displays the updated status.

```
naviseccli -h ss1_spa mirror -enablepath ss27_spa -connectiontype  
iscsi
```

```
naviseccli -h ss1_spa mirror -async -info
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`.

mirror -async -fractureimage

Fractures an image in a mirror.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to start an administrative fracture of a secondary image. It also prevents scheduled updates from starting in the future. You must direct the command to the storage system that holds the primary image.

(An automatic fracture, called a system fracture, occurs if a secondary image becomes inaccessible to the primary image. After a system fracture, the system periodically tries to communicate with the secondary image and — if it succeeds — starts synchronization based on the fracture log. If you issue the `-fractureimage` command while the mirror is system fractured, the system will not try to communicate or resynchronize until you issue a `-syncimage` command.)

The software displays a message to confirm the fracture:

```
Do you want to fracture the image image UID (y/n)?
```

To fracture the mirror, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -async -fractureimage -name name | -mirroruid uid
-imageuid uid
[-o]
```

OPTIONS

`-name name`

Specifies the name you gave to the mirror with the `-create` command.

`-mirroruid uid`

Specifies the unique 16-byte mirror ID (World Wide Name) that the software assigned when it created the mirror.

`-imageuid uid`

Specifies the unique 8-byte image ID of the secondary image to be fractured.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa` of the mirror named `db_mirror`, this command fractures the secondary image with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -fractureimage -name  
db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

```
Do you really want to fracture the image(y/n)? y
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -info

Displays general information about asynchronous mirroring in a storage system.

PREREQUISITES

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

DESCRIPTION

This command, when used with `naviseccli` and no other switches, lists all general information for asynchronous mirror support. Issuing this command with one or more options displays the information based on the option issued. See the example for sample output.

SYNTAX

```
mirror -async -info [-cancreate]
[-maxmirrors]
[-maxremoteimages]
[-mirroredluns]
[-mirrorableluns]
[-systems [-connectiontype]]
```

OPTIONS

-cancreate

Indicates whether an asynchronous mirror can be created in this storage system: yes or no.

-maxmirrors

Displays the maximum number of mirrors that can be created in this storage system.

-maxremoteimages

Specifies the maximum number of secondary images allowed. A maximum of one image is allowed.

-mirroredluns

Lists the LUN numbers of LUNs that are mirrored in this storage system. The list includes LUNs used as both secondary and primary images.

-mirrorableluns

Lists LUNs that can be mirrored in this storage system but are not mirrored at present.

-systems

Lists remote storage systems that can be used for a secondary image.

-connectiontype

Displays the connection type (fibre or iSCSI) for the enabled mirror. For unexpected error cases, it displays the value, unknown.

-logs

Lists LUNs that have been assigned for the write intent log on this storage system.

EXAMPLE

For SP `ss1_spa`, this command displays potential and current remote mirror information.

```
naviseccli -h ss1_spa mirror -async -info
```

```
Can a mirror be created on this system:YES
```

```
Maximum number of remote images: 1 Logical
```

```
Units that are mirrored in Primary Images:128 5
```

```
Logical Units that are mirrored in Secondary Images:0
```

```
Logical Units that can be mirrored:1 2 3 4 10 11 12 13  
14 21 22 58 59
```

```
Remote systems that can be enabled for mirroring:
```

```
Array UID: 50:06:01:60:90:20:27:47
```

```
Status: Enabled on both SPs Enabled Connection
```

```
Type(s): iSCSI Disabled Connection Type(s): Fibre Maximum
```

```
number of possible Mirrors:100
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a Command is not supported error message is printed to stdout.

mirror -async -list

Displays information about existing asynchronous mirrors.

PREREQUISITES

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

DESCRIPTION

The `naviseccli -async mirror` command with `-list` and no other switches lists all information for all mirrors on the storage system. Issuing this command with one or more options displays the information based on the option(s) issued. You can direct the command to any storage system; however, only the storage system holding the primary image of a mirror will always have the latest information about the mirror. Status obtained from other storage systems may be stale; do not use it unless the primary storage system is unavailable. See the example for sample output.

SYNTAX

```
mirror -async -list [-name [[name]] [-mirroruid] [uid]
[-description]
[-faulted]
[-images]
[-lun]
[-ownergroupuid]
[-ownergroupname]
[-requiredimages]
[-state]
[-timestamp]
[-transition]
```

OPTIONS

`-name` *name*

Specifies the name you gave the mirror with the `-create` function. You can use the mirror name or unique ID to obtain information on a mirror.

`-mirroruid` *uid*

Specifies the unique 16-byte remote mirror ID (World Wide Name) the software assigned when it created the mirror.

`-description`

Lists the meaningful title for the mirror, if any. The software does not use the description to identify the mirror.

-faulted

Lists whether the remote mirror is faulted: yes or no. If it is faulted, the mirror is either fractured or in the attention state.

-images

Displays image information, such as the role of the image (primary or secondary), image ID, WWN of the LUN containing the image, recovery policy, sync rate, image state, image condition, and image transitioning. It also displays the update type (Manual, Time elapsed since the start of last update, and Time elapsed since the end of the last update), the update period in minutes (displays N/A where not applicable), the time in minutes since the last update, the time in minutes until the next update, the last image error (why the image is faulted; displays the Not Available if no error).

-imagecount

Lists the current number of images that compose the mirror. This count includes the primary image.

-imagesize

Lists the size of the image (capacity of the LUN containing the image) in user blocks.

-lun

Lists the LUN number of the local image.

-ownergroupuid

Lists the unique ID value of the consistency group of which the mirror is a member.

-ownergroupname

Lists the unique name of the consistency group of which the mirror is a member.

-requiredimages

Lists the minimum number of secondary images required for the mirror.

-state

Describes the operational state of the mirror. Values are: active and attention.

-timestamp

Returns the clock time when the last update started.

-transition

Specifies whether the mirror is in a changing state.

EXAMPLE

```
naviseccli -h ssl_spa mirror -async -list
```

```
MirrorView Name:  pbc async ml 80
                  MirrorView Description:
                  MirrorView UID:
54:33:20:90:60:01:06:50:03:00:00:00:00:00:00:00
                  Logical Unit Numbers:  80
                  Remote Mirror Status:  Mirrored
                  MirrorView State:  Active
                  MirrorView Faulted:  NO
                  MirrorView Transitioning:  NO
                  Minimum number of images required:  0
                  Image Size:  409600
                  Image Count:  2
                  Images:
                  Image UID:  50:06:01:60:90:20:33:54
                  Is Image Primary:  YES
                  Logical Unit UID:
60:06:01:60:2F:E7:0C:00:3C:5B:04:A2:50:4C:D8:11
                  Image Condition:  Unknown Image Condition

Image UID:  50:06:01:60:90:20:29:85
Is Image Primary:  NO
Logical Unit UID:  60:06:01:60:3F:E1:0D:00:A4:C9:01:63:B6:4C:D8:11
Image State:  Synchronized
Image Condition:  Normal
Recovery Policy:  Automatic
Synchronization Rate:  Medium
Image Faulted:  NO
Image Transitioning:  NO
Synchronizing Progress(%):  0
Update Type:  Manual
Time in secs since previous update:  12550
Time in secs until next update:  Not Available
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a Command is not supported error message is printed to stdout. You may also receive other error messages.

mirror -async -promoteimage

Promotes a secondary image to primary.

PREREQUISITES

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

DESCRIPTION

This command promotes a secondary image to the role of the primary image. The command is successful if it can promote the secondary to a primary that is server-accessible. You must direct this command to a storage system that holds the secondary image — never to the storage system that holds the primary. The command will fail if you issue it to the storage system that holds the primary image.

You can promote a secondary image if the secondary is in a synchronized or consistent state. There must be no current update, in addition to the state restrictions listed.

When you promote a secondary image to primary, the software assigns a new mirror UID to the promoted image to distinguish it from the old mirror. The new status of the old primary image depends on whether the old primary image is accessible when promotion occurs.

After promoting an image, you must explicitly add the newly promoted image to an appropriate storage group. EMC recommends that you also remove the old primary from its storage group, if possible.

In normal operation when there is no failure, you can (but probably would not want to) promote a secondary image. If you do so, the primary image is demoted to secondary. The new primary image inherits the mirror properties of the demoted primary image. Before issuing this command consider the following:

- ◆ If possible, the secondary image should be synchronized (synchronized state) with the primary image. If the secondary image is in the consistent state (not synchronized state) when you promote it to primary image, data will probably be lost during the promotion process. If possible, the secondary image should be synchronized (synchronized state) with the primary image. If the secondary image is in the consistent state (not synchronized state) when you promote it to primary image, data will probably be lost during the promotion process.

The secondary image to be promoted must be in a consistent or synchronized state. If the secondary image is out-of-sync or synchronizing, then you cannot promote it; an error will result.

- ◆ You must stop I/O to an accessible primary image before you issue this command. If I/O continues to the original primary, data may be lost as the software demotes the primary to secondary and/or an SP may reboot unexpectedly during the attempted promotion.

SYNTAX

```
mirror -async -promoteimage -name name | -mirroruid uid
-imageuid imageuid
[-type normal|oos|local ] [-o]
[-type type]
[-o]
```

OPTIONS

-name *name*

Specifies the name you gave the mirror with the **-create** command. You can use the mirror name or unique ID to promote a secondary image.

-mirroruid *uid*

Specifies the unique 16-byte mirror ID (World Wide Name) that was assigned when the mirror was created. For any letters in the ID, see comments under **-imageuid** below.

-imageuid *uid*

Specifies the unique 8-byte image ID of the secondary image you want to promote. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the **mirror -list** function).

-type normal|oos|local

Checks the status of the mirror images. Normal (the default) directs the driver to perform internal safety checks and return the appropriate error if the promotion would cause an out-of-sync state for the old primary or remote image if the mirror is unreachable. OOS indicates that no safety check is performed, and the promote will occur even if it results in an out-of-sync state for the old primary. Local results in promoting the local image only to the primary. The driver attempts to contact the old primary, and if it reaches it, the old primary removes its secondary image.

-o

Runs the command without prompting for confirmation.

EXAMPLE

For SP **ss1_spa**, the owner of the remote mirror named **db_mirror**, this command promotes the secondary image with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -promoteimage
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -removeimage

Removes a secondary image from a mirror.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to remove a secondary image from the remote mirror. You must issue this command to a storage system that holds the primary image.

Note: If an update is in progress, this command will not execute. We recommend that you fracture the secondary image first, and then you can issue the `mirror -async -removeimage` command.

If you omit the `-override` option, the software displays a message to confirm the remove operation:

```
Removing of the image will stop mirroring of data to the specified
image. It may cause the minimum number of images required parameter
to be violated that will cause the mirror to go into attention state.
Do you really want to remove the image(y/n)?
```

Answering yes will remove the image.

SYNTAX

```
mirror -async -removeimage -name name | -mirroruid mirroruid
-arrayhost sp-hostname | sp-IP-address | -imageuid imageuid [-o]
```

OPTIONS

`-name name`

Specifies the name you gave the mirror with the `-create` command. You can use the mirror name or unique ID to remove a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

`-arrayhost sp-hostname | sp-IP-address`

Specifies the hostname or network address of the remote SP that holds the secondary image you want to remove. If you specify this switch, the driver(s) from the secondary LUN stack is removed.

-imageuid *image-uid*

Specifies the unique 8-byte image ID of the secondary image you want to remove. For any letters in the ID, you must specify the same case as the software assigned to the image when it created the mirror (use the **mirror -list** command). If you specify this switch, you must explicitly remove the driver(s) from the secondary LUN stack using the **-setfeature** command.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP address **ss1_spa**, the owner of the mirror named **db_mirror**, this command removes the image with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -removeimage  
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a **Command is not supported** error message is printed to **stdout**. You may also receive other error messages.

mirror -async-sync -setfeature

Adds or removes an asynchronous mirror driver to or from the LUN stack.

PREREQUISITES

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

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 32](#)) to add the asynchronous mirror driver to or remove it from the LUN stack depending on whether you specify the `-on` or `-off` option. You must issue this command to the storage system where the LUN resides.

Note: `-setfeature` is needed only if you want to use the `-arrayuid` switch to specify the remote storage system for the `-addimage` command.

Note: When you add a synchronous mirror driver to the LUN stack using the `-on` switch, MirrorView/S automatically allocates and retains additional memory. This additional memory counts against the total replication software memory budget.

SYNTAX

```
mirror -async -setfeature -on | -off -lun lun_number | -lunuid uid
```

OPTIONS

`-on`

Adds the driver to the LUN stack.

`-off`

Removes the driver from the LUN stack.

`-lun lun_number`

Specifies the LUN number to which you want to add or from which you need to remove the driver.

`-lunuid uid`

Specifies the unique 16-byte identification number (World Wide Name) of the LUN to which you will add the driver or from which you will remove the driver.

EXAMPLE

For SP **ss1_spa**, this command starts adding the driver to LUN 112.

```
naviseccli -h ss1_spa mirror -async -setfeature -on  
-lun 112
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -syncimage

Starts an update of the secondary image.

PREREQUISITES

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

DESCRIPTION

Use this command with `naviseccli` (see [naviseccli on page 32](#)) to begin the synchronization operation on a secondary image that was previously administratively fractured or on one where the mirror image condition is normal. You must issue this command to the storage system that holds the primary image. The secondary image should be fractured when performing this operation. If you specify a secondary image that is not fractured, an internal error message will be returned.

The software displays a confirmation message:

```
During synchronization, array performance will be impacted. Do you
want to sync the image (y/n)?
```

Answering `y` will start synchronization.

SYNTAX

```
mirror -async -syncimage -name name | -mirroruid mirroruid -imageuid imageuid
[-o]
```

OPTIONS

`-name name`

Specifies the name given to the mirror with the `-create` command. You can use the mirror name or unique ID to update a secondary image.

`-mirroruid uid`

Specifies the unique 16-byte remote mirror ID (World Wide Name) assigned when the software created the mirror. For any letters in the ID, see comments under `-imageuid` below.

`-imageuid imageuid`

Specifies the unique 8-byte image ID of the secondary image you will synchronize. For any letters in the ID, you must specify the same case as the software assigned to the image when the software created the mirror (use the `mirror -list` command).

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, the owner of the mirror named `db_mirror`, this command starts synchronizing the image with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -syncimage  
-name db_mirror -imageuid 50:06:06:10:00:FD:A1:6
```

During synchronization, array performance will be impacted. Do you want to sync the image (y/n)? `y`

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

MirrorView/A Consistency Group Commands

This chapter explains the commands for creating and managing groups of mirrors on CX3 series, CX series, and AX4-5 series storage systems with the EMC MirrorView/A option. This manual refers to groups of mirrors as *consistency groups*.

Note: The LUNs in a MirrorView/A consistency group must reside on the same storage system, unlike the volumes in a Symmetrix consistency group, which can reside on multiple storage systems.

Major topics are:

- ◆ [mirror -async -creategroup on page 74](#)
- ◆ [mirror -async -destroygroup on page 77](#)
- ◆ [mirror -async -addtogroup on page 79](#)
- ◆ [mirror -async -removefromgroup on page 81](#)
- ◆ [mirror -async -change group on page 83](#)
- ◆ [mirror -async -syncgroup on page 85](#)
- ◆ [mirror -async -fracturegroup on page 87](#)
- ◆ [mirror -async -promotegroup on page 89](#)
- ◆ [mirror -async -listgroups on page 91](#)

mirror -async -creategroup

Establishes a consistency group to which you can later add mirrors.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to create a consistency group to which you can add mirrors with subsequent commands. You must direct the command to the storage system on which you are creating the MirrorView/A group.

Note: On AX4-5, CX500, CX3 model 10, and CX3 model 20 systems, you can have 8 consistency groups, and each consistency group can have up to 8 mirrors. On CX700, CX3 model 40, and CX3 model 80 systems, you can have 16 consistency groups, and each consistency group can have up to 16 mirrors. On CX4 storage systems, you can have 64 consistency groups. On CX4-120 and CX4-240 storage systems, each consistency group can have up to 32 mirrors. On CX4-480 and CX4-960 storage systems, each consistency group can have up to 64 mirrors.

SYNTAX

```
mirror -async -creategroup -name name | -mirroruid uid
[-description description]
[-manualupdate | -delaystart minutes | -enddelay minutes]
[-syncrate rate] [-o]
[-recoverypolicy policy] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group. Use only alphanumeric characters for the consistency group name. The name can be up to 32 characters.

-description *description*

Specifies a textual description of the consistency group.

-manualupdate

Specifies that you must explicitly start each update of the group. If you specify neither **-manualupdate**, **-delaystart**, nor **-enddelay**, the default is **-enddelay 60**.

Note: When a consistency group begins an update, the mirrors within the group begin updating and end whenever each individual update finishes. The consistency group update completes when the last mirror completes. If an update is not running, you can manually start an update even if an automatic update interval is set.

-delaystart *minutes*

Specifies the time (in minutes) from the start of the previous update until the start of the next one. If the previous update has not completed, the next update is delayed until the update in progress has completed. If you specify neither **-manualupdate**, **-delaystart**, nor **-enddelay**, the default is **-enddelay 60**.

-enddelay *minutes*

Specifies the time (in minutes) from the end of the previous update until the start of the next one. Specifying a value of 0 results in the group being updated as fast as possible. If you specify neither **-manualupdate**, **-delaystart**, nor **-enddelay**, the default is **-enddelay 60**.

-syncrate *rate*

Specifies a relative value (low, medium, or high) for the priority of completing updates. Values for rate are low, medium, or high. High completes updates faster, but may significantly affect storage-system performance for server I/O requests. Low completes updates slower, but also minimizes the impact on other storage-system operations.

-recoverypolicy *policy*

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

manual: Update of the secondary image does not begin until you explicitly issue a **synchronize** command.

auto: Update of the secondary image automatically begins as soon as the connection between the two storage systems is again operational.

-o

Executes the command without prompting for confirmation.

Note: All the switches, except **-name** and **-o**, are meaningful only if the storage system receiving the command hosts the primary images for the mirrors in the consistency group.

EXAMPLE

For SP `ss1_spa`, this command creates a consistency group with the name **dbgroup**.

```
naviseccli -h ss1_spa mirror -async -creategroup -name dbgroup
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -destroygroup

Destroys a consistency group.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to destroy a consistency group. It is the counterpart of `-creategroup`.

Note: Make sure that the consistency group is empty before destroying it.

If you omit the override option, the system displays a message to confirm the destruction of the consistency group:

```
Do you really want to destroy the group <group-name>? (y/n)?
```

To destroy the consistency group, answer `y`; to take no action, enter `n`.

SYNTAX

```
mirror -async -destroygroup -name name | -groupid id
[-force] [-o]
```

OPTIONS

-name *name*

Specifies the name you gave the mirror with the `-creategroup` command. You can also use the consistency group ID.

-groupid *id*

Specifies the consistency group unique ID assigned when the consistency group was created.

-force

Destroys the consistency group even if it is not empty.

Note: Particularly in a failure situation you can use the `-force` switch to destroy a group that still has members.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ssl_spa`, this command destroys the consistency group that was created with the name `dbgroup`.

```
naviseccli -h ssl_spa mirror -async -destroygroup dbgroup  
Do you really want to destroy the group dbgroup?(y/n)? y
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages if the command fails.

mirror -async -addtogroup

Adds a mirror to a previously created consistency group.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to add one mirror at a time to a consistency group.

Note: The mirror you specify must have its primary image on the storage system. No two secondary images of mirrors in the consistency group can be on different storage systems.

SYNTAX

```
mirror -async -addtogroup -name name | -groupid id
-mirrorname name | -mirroruid id
```

OPTIONS

-name *name*

Specifies the name of the consistency group to which you want to add the mirror.

-groupid *id*

Specifies the unique ID of the consistency group to which you want to add the mirror.

-mirrorname *name*

Specifies the name of the mirror you want to add to the consistency group.

-mirroruid *id*

Specifies the World Wide Name of the mirror you want to add to the consistency group.

EXAMPLE

For SP `ss1_spa`, this command adds the mirror, **Sales Accounts Mirror to Europe**, to the consistency group, **dbgroup**.

```
naviseccli -h ss1_spa mirror -async -addtogroup
-name dbgroup -mirrorname "Sales Accounts Mirror to Europe"
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -removefromgroup

Removes a mirror from a consistency group.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` (see [naviseccli on page 32](#)) to remove one mirror at a time from the consistency group.

Note: The mirror you specify must have its primary image on the storage system. This does not apply if you use the `-force` switch.

If you omit the `-o` option, the software displays a message to confirm the remove operation. Answering yes will remove the mirror.

SYNTAX

```
mirror -async -removefromgroup -name name | -groupid id
-mirrorname name | -mirroruid id[-force] [-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group from which you want to remove the mirror.

`-groupid id`

Specifies the unique ID of the consistency group to which you want to add the mirror.

`-mirrorname name`

Specifies the name of the mirror you want to remove from the consistency group.

`-mirroruid id`

Specifies the World Wide Name of the mirror you want to remove from the consistency group.

`-force`

Removes the mirror even if the connection with the remote storage system is not functioning.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, the owner of the consistency group named `dbgroup`, this command removes the mirror with the unique ID shown.

```
naviseccli -h ss1_spa mirror -async -removefromgroup  
-name dbgroup -mirroruid
```

```
83:5A:20:90:60:01:06:50:01:00:00:00:00:00:00:00
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -change group

Changes consistency group properties.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to change one or more parameters of a consistency group.

You must specify at least one of the optional switches. If you omit the override option, then it displays a message similar to the following for each parameter to confirm the change.

```
Do you want to change the group synchronization rate from Medium to
High (y/n)?
```

To change a parameter, enter `y`.

SYNTAX

```
mirror -async -change group -name name | -groupid id
[-description description]
[-manualupdate | -delaystart minutes | -enddelay minutes]
[-syncrate rate]
[-recoverypolicy policy] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to change. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to change.

-description *description*

Specifies a textual description for the consistency group.

-manualupdate

Specifies that you must explicitly start each update of the group.

-delaystart *minutes*

Specifies the time (in minutes) from the start of the previous update until the start of the next one. If the previous update has not completed, the next update is delayed until the update in progress has completed.

-enddelay *minutes*

Specifies the time (in minutes) from the end of the previous update until the start of the next one. Specifying a value of 0 results in the group being updated as fast as possible.

-syncrate *rate*

Specifies a relative value (low, medium, or high) for the priority of completing updates. Values for rate are low, medium, or high. High completes updates faster, but may significantly affect storage-system performance for server I/O requests. Low completes updates slower, but also minimizes the impact on other storage-system operations.

-recoverypolicy *policy*

Specifies the policy for recovering the mirrors in a group after a failure causes a system fracture. Values are:

manual: Update of the secondary image does not begin until you explicitly issue a **synchronize** command.

auto: Update of the secondary image automatically begins as soon as the connection between the two storage systems is again operational.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For storage system `server1`, consistency group `dbgroup`, this command changes the synchronization rate to `high`.

```
naviseccli -h server1 mirror -async -change group -name
dbgroup -syncrate high
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

Use the **-listgroups** command on the mirror to confirm that the changes were made.

mirror -async -syncgroup

Starts an update of the consistency group.

PREREQUISITES

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

DESCRIPTION

This command is used with `naviseccli` to begin the synchronization operation on a consistency group. All the images in the consistency group must be primary.

SYNTAX

```
mirror -async -syncgroup -name name | -groupid id
[-synctype startnow]
[-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to update or synchronize. You can specify either the consistency group name or id (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to update or synchronize.

-synctype *startnow*

Starts synchronization immediately. *startnow* is the only currently supported option.

-o

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, the owner of the consistency group named `dbgroup`, this command starts synchronizing the consistency group immediately.

```
naviseccli -h ss1_spa mirror -async -syncgroup -name dbgroup
-synctype startnow
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -fracturegroup

Fractures all image(s) in a consistency group.

PREREQUISITES

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

DESCRIPTION

The `naviseccli mirror -async` command with `-fracturegroup` starts an administrative fracture of all images in the consistency group, even if they are already system fractured. The consistency group cannot be in a scrambled state, and all images must be primary. If the consistency group is in the process of updating, it stops.

SYNTAX

```
mirror -async -fracturegroup -name name | -groupid id
[-fracturetype abort]
[-o]
```

OPTIONS

`-name name`

Specifies the name of the consistency group you want to fracture. You can specify either the consistency group name or ID (next).

`-groupid id`

Specifies the unique ID of the consistency group you want to fracture.

`-fracturetype abort`

Specifies that you want to abort an update if one is in progress. The secondary retains any extents that have been copied to the secondary, and therefore do not need to be resent on the next synchronization.

`-o`

Executes the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, with the consistency group named `dbgroup`, this command fractures all the images in the consistency group.

```
naviseccli -h ssl_spa mirror -async -fracturegroup  
-name dbggroup
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a Command is not supported error message is printed to stdout. You may also receive other error messages.

mirror -async -promotegroup

Promotes all secondary images in a consistency group to primary images.

PREREQUISITES

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

DESCRIPTION

This command promotes all secondary images to the role of the primary images. This command is used with `naviseccli` (see [naviseccli on page 32](#)).

SYNTAX

```
mirror -async -promotegroup -name name | -groupid id
[-type type] [-o]
```

OPTIONS

-name *name*

Specifies the name of the consistency group you want to promote. You can specify either the consistency group name or ID (next).

-groupid *id*

Specifies the unique ID of the consistency group you want to promote.

-type *type*

Specifies the type of promotion. Values are:

`normal` (default), which causes a failure if the connection to the primary storage system is not functioning or if at least one set of the images to promote are not synchronized with the primaries. Otherwise, it will destroy the original mirrors and create new ones with the targeted secondaries as primaries and the original primaries as secondaries.

`local`, which does not attempt to add back the original primaries as secondaries, but always creates new mirrors with only primary images consisting of the targeted secondaries.

`oos` (out-of-sync), which attempts a normal promotion, reversing the original roles of primaries and secondaries even if they are not all synchronized. If the connection to the primary storage system is not functioning, a local promotion will occur.

-o

Runs the command without prompting for confirmation.

EXAMPLE

For SP `ss1_spa`, the owner of the consistency group named `dbgroup`, this command promotes all secondary images:

```
naviseccli -h ss1_spa mirror -async -promotegroup -name  
dbgroup
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

mirror -async -listgroups

Displays information about existing consistency groups.

PREREQUISITES

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

DESCRIPTION

The `naviseccli mirror -async` command with `-listgroups` and no other switches lists all information for all asynchronous consistency groups on the storage system. Issuing this command with one or more options displays the information based on the option(s) issued. This command is used with `naviseccli` (see [naviseccli on page 32](#)).

SYNTAX

```
mirror -async -listgroups [-name[name]] [-groupid[-id]]
[-feature] [-update] [-nextupdate]
[-syncrate] [-state] [-role] [-mirrors]
[-condition] [-recoverypolicy] [-description]
```

OPTIONS

-name *name*

Specifies the name of the consistency group whose information you want to display. If you do not specify the name of the consistency group with this switch, that piece of information will be provided for the consistency group(s) being displayed.

-groupid *id*

Specifies the unique ID of the consistency group whose information you want to display.

-feature

Displays the maximum number of consistency groups allowed on the storage system and the maximum number of members per consistency group.

-update

Displays the type of update as either manual, Time elapsed since the start of the last update (in minutes), or Time elapsed since the end of the last update (in minutes). On the next line it also displays the Update Period in minutes. If the Update Period is not applicable, it displays N/A.

-nextupdate

Displays the number of minutes until the next update.

-syncrate

Displays the rate at which data is transferred during updates and initial synchronizations.

-state

Displays Synchronized, Out-of-Sync, Synchronizing, Consistent, Rollingback, Quasi-Consistent, Scrambled, Incomplete, Local Only, or Empty.

-role

Displays Primary, Secondary, or Unknown.

-mirrors

Displays the mirror name, mirror uid, the IDs of both images in the pair, and the storage system ID of the paired image.

-condition

Displays the consistency group condition as Normal, Initializing, Updating, Admin Fractured, System Fractured, Waiting on Admin.

-recoverypolicy

Displays either manual or automatic to describe whether the system will wait for you to restart updates in the event of a system fracture, after the problem is corrected.

-description

Displays the description associated with the consistency group.

EXAMPLE

```
naviseccli -h 10.14.5.200 mirror -async -listgroups
```

```
Group Name: taraFARGroup
```

```
Group ID: 50:06:01:60:90:20:5A:83:00:00:00:00
```

```
Description:
```

```
Maximum Number of Groups Allowed: 16
```

```
Maximum Number of Mirrors per Group: 16
```

```
Update Type: Manual
```

```
Time to next update: 0
```

```
State: Quasi Consistent
```

```
Role: Secondary
```

```
Mirror Name: taraFAR202
Mirror WWN: C6:05:20:88:60:01:06:50:08:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Mirror Name: taraFAR201
Mirror WWN: C6:05:20:88:60:01:06:50:07:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Mirror Name: taraFAR200
Mirror WWN: C6:05:20:88:60:01:06:50:06:00:00:00:00:00:00
Primary WWN: 50:06:01:60:88:20:05:C6
Secondary WWN: 50:06:01:60:90:20:5A:83
Array hosting remote image: 50:06:01:60:88:20:05:C6
Condition: Admin Fractured - Intervention required
Sync Rate: Medium
Recovery Policy: Manual
Group Name: HWLocalPromoteTestGroup
Group ID: 50:06:01:60:90:20:5A:83:01:00:00:00
Description:
Maximum Number of Groups Allowed: 16
Maximum Number of Mirrors per Group: 16
Update Type: Manual
Time to next update: 0
State: Synchronized
Role: Secondary
Mirror Name: HWLocalPromoteTestMirror
Mirror WWN: C6:05:20:88:60:01:06:50:09:00:00:00:00:00:00
```

```
Primary WWN: 50:06:01:60:88:20:05:C6  
Secondary WWN: 50:06:01:60:90:20:5A:83  
Array hosting remote image: 50:06:01:60:88:20:05:C6  
Condition: Normal  
Sync Rate: Medium  
Recovery Policy: Manual
```

OUTPUT

If the version of FLARE software running on the SP does not support this command, a `Command is not supported` error message is printed to `stdout`. You may also receive other error messages.

Appendix A

Responding to Failures

This appendix describes how MirrorView/A and consistency groups handle failures.

Topics include:

- ◆ [How MirrorView/A handles failures on page 96](#)
- ◆ [How consistency groups handle failures on page 105](#)
- ◆ [iSCSI-related failure conditions on page 106](#)

How MirrorView/A handles failures

When a failure occurs during normal operations, MirrorView/A lets you perform several actions to recover.

In recovering from failures, MirrorView/A achieves two goals:

- ◆ Preserves data integrity
- ◆ Minimizes the amount of time that data is unavailable to the user

Access to the primary SP fails

If an SP that owns mirrored LUNs on the primary system fails, MirrorView/A on the other SP takes ownership of those mirrored LUNs by trespassing them when something on the server (like PowerPath) initiates the trespass. This allows mirroring to continue, provided the server is set up properly to handle the failover (for example, a Windows server with PowerPath). When the primary LUN is trespassed, MirrorView/A sends a trespass request to any secondary images when the next update starts. Therefore, you may notice that the mirrored LUNs on the secondary system have moved from SP A to SP B, or vice versa. MirrorView/A keeps the SP ownership the same on the primary and secondary systems during updates. If the primary image is on SP A, then the secondary image will be on SP A. This may not occur until the start of the next update.

Primary image fails

If the storage system controlling the primary image fails, access to the mirror stops until you either repair the storage system or promote a secondary image of the mirror to primary.

Promoting a secondary image to primary

In situations where you must replace the primary storage system due to a catastrophic failure, you can use a promotion to access data on the secondary storage system. To recover and restore I/O access, you must promote a secondary mirror image to the role of primary mirror image, so that a server can access it.

Note: You can also promote a secondary image even if there has not been a catastrophic failure.

If the primary image and secondary image can communicate with each other, then when the secondary image is promoted, the former primary image is demoted to a secondary image.

To promote a secondary image, the following conditions must be true:

- ◆ You must direct the `naviseccli mirror` commands to the storage system holding the secondary image.
- ◆ The state of the secondary image you will promote must be either Consistent or Synchronized.
- ◆ An update is not currently transferring data for this mirror.



CAUTION Promoting a secondary image will cause loss of data written to the primary image after the start of the last completed update. If any updates have been made to the primary image since that time, a full resynchronization of the mirror will be required after the promotion. Also, if an update is currently active (that is, transferring data), the promotion will not be allowed; allow the update to complete and the image to transition into the Synchronized state; then perform the promotion. An alternative to allowing the update to complete is to fracture the image.

In a failure situation, before promoting a secondary image to a primary image:

- ◆ If the existing primary image is accessible, remove the primary image from any storage groups before promoting the secondary image to avoid I/O and therefore inconsistent data.
- ◆ Ensure that no I/O, either generated from a server or by an update in progress, is occurring in the asynchronous mirror.
- ◆ If the existing primary is available, make sure that it lists the secondary image that is to be promoted as "synchronized."

To promote a secondary image to a primary image:

1. Issue the `mirror -async -promoteimage` command. See [mirror -async -promoteimage on page 63](#)

Note: If the original primary storage system failed, remove the primary storage system from the domain.

If you do not specify the `-type` switch, the command performs a normal promote.

2. Add the newly promoted image to a storage group if necessary.

At some point later, you can also perform the following steps:

1. Verify that the failed storage system is not the master of the domain.

If it is, assign another storage system to be the master. See the *EMC Navisphere Command Line Interface (CLI) Reference*.

2. Verify that the failed storage system is not a portal.

If it is a portal, remove the portal and configure a different storage system as a portal. See the *EMC Navisphere Command Line Interface (CLI) Reference*.

Example promote mirror scenario #1

You attempt to promote a mirror that has a secondary image, but the connection between the storage systems is not working. The secondary image indicates that it is synchronized, when it is actually system fractured and consistent. An error, Existing image unreachable, appears. You can investigate the reason for the loss of connectivity and correct the problem before continuing with the promotion, or you can select the Local Only Promote option to complete the promotion.

If you select Local Only Promote, the software promotes the local mirror and attempts to contact the original primary image and remove the promoted image from the mirror. In the case described here, it cannot contact the other storage system, so it converts the local image to a primary image in a mirror with no secondary images.

Note: In this scenario, a Force Promote has exactly the same effect as the Local Only Promote operation.

Since the software cannot contact the remote storage system, the original mirror still exists on the storage system originally hosting the primary image. However, even if connectivity is restored, any attempt to start an update will fail (since the secondary has been promoted), and the secondary image will remain administratively fractured forever. You should use **Force Destroy** to remove this original mirror.

Example promote mirror scenario #2

You attempt to promote a mirror whose secondary image is in the Consistent state. An error, Existing primary will be out-of-sync, appears. If possible, allow the secondary to go to the Synchronized state (for example, stop application I/O to the primary image, flush data from the server, start an update and wait for it to complete). You can then promote the secondary without requiring a full resynchronization. Otherwise, you can select either the Force Promote or the Local Only Promote option to continue the promotion. In either case, you must perform a full resynchronization before you have the mirror providing protection for your data again.

Failure of the secondary image

When a primary image cannot communicate with a secondary image, it marks the secondary as unreachable and will system or administratively fracture the secondary image. Verify the Last Image Error field for details about the cause of the failure and take the appropriate action to fix it. The secondary image is marked System Fractured. The loss of communication may be due to a link between storage systems, an SP failure on the secondary storage system, or some other failure on the secondary storage system. In the event of the communication failure, the secondary image remains a member of the mirror.

If the mirror is set for automatic recovery, an update automatically starts once the secondary storage system is again accessible. Otherwise, you must manually start the update.

Promoting a secondary image when there is no failure

You may want to promote your secondary image even if no failure occurs on the storage systems. For example, you may want to test your disaster recovery procedure before a real failure occurs, or perhaps the server attached to the primary storage system failed, and you must resume operations using the server attached to the secondary storage system.

If the original primary is accessible when you promote the secondary, the software verifies whether the images are identical. If possible, the secondary image should be in the Synchronized state (stop application I/O, flush data from the servers, start and update and wait for it to complete). If the images are identical, they swap roles, resynchronization is not necessary, and the promotion is complete. If the images are potentially different (that is, the secondary image is not in the Synchronized state), then you must specify the type of promotion. As part of a promotion, any secondary images in the new mirror are set to manual recovery.

Mirror before promotion	Mirror after promotion
mirror ID = aaa	mirror ID = bbb
primary image = LUN xxxx	primary image = LUN yyyy
secondary image = LUN yyyy	secondary image = LUN xxxx

If the images are not synchronized, you can choose to forcibly promote (oos), locally promote, or to not promote. OOS promote and local promote require a full resynchronization of the data before mirrored protection is again in effect.

Summary of MirrorView/A failures

[Table 4 on page 99](#) shows how MirrorView/A might help you recover from system failure at the primary and secondary sites. It assumes that the secondary image of the mirror is in either the Synchronized or Consistent state.

Table 4. Summary of MirrorView/A failures

Event	Result and recovery
Loss of access to primary image LUN	Check connections between the server and storage system, including zoning and correct operation of any switches. Check for an SP reboot.

Table 4. Summary of MirrorView/A failures *(continued)*

Event	Result and recovery
Secondary SP is rebooted	If the secondary SP reboots, for example, due to a software failure, an explicit command or the SP is physically pulled and reseated, you may see the secondary image become system fractured. It is also possible for the secondary to become administratively fractured, in which case simply synchronize the image.
Server accessing primary image fails	Catastrophic failure. I/O stops. After two more update periods, if the secondary image is in the Consistent state, it will transition to the Synchronized state. Nothing more happens until the server is repaired or replaced, or a secondary image is promoted.

Table 4. Summary of MirrorView/A failures *(continued)*

Event	Result and recovery
Storage system running primary image fails	<p>Option 1 - Catastrophic failure. The mirror is left in the state it was already in. If the secondary image is in either the Consistent or Synchronized state, it may be promoted to provide access to your data.</p> <hr/> <p>Note: Any writes to the primary since the start of the last completed update are not propagated to the secondary.</p> <hr/>
	<p>Option 2 - Noncatastrophic failure, repair is feasible.</p> <p>The administrator has the problem fixed, and then normal production operation can resume.</p>
	<p>Option 3 - Only one SP fails.</p> <p>If the SP that controls the mirror fails, software on the server (for example, PowerPath) can detect the failure. This software can then cause control of the mirror to be transferred to the surviving SP, and normal operations can continue. If such software is not running on the server, then you must either manually transfer control using Navisphere, or access to the mirror stops until the SP is back in service.</p> <p>If the SP that does not control the LUN fails, then mirroring continues as normal.</p>

Table 4. Summary of MirrorView/A failures *(continued)*

Event	Result and recovery
Storage system running secondary image fails	<ul style="list-style-type: none"> ◆ If the SP that does not control the secondary image fails, nothing happens with respect to this mirror. ◆ If the SP that controls the mirror fails (or both SPs fail or a catastrophic failure of the entire system occurs), the secondary image of the mirror will become system fractured at the start of the next update. If the mirror is appropriately configured, this will cause the mirror to go into the Attention state, although I/O will continue as normal to the primary image. <p>The administrator has a choice: If the secondary can easily be fixed (for example, if someone pulled out a cable), then the administrator can have it fixed and let things resume. Otherwise, to regain protection of your data and ensure that you have another storage system available, you can force destroy the existing mirror, recreate it, and add a secondary image on another working storage system. Protection is not established until the full synchronization of the secondary image completes.</p>
Loss of connection between storage systems (indicated by system fractures)	<ul style="list-style-type: none"> ◆ Verify zoning of ports and switches for the MirrorView connection between storage systems. ◆ Verify all cables are properly seated and in the correct location. ◆ Check the switch logs to ensure no connection or firmware issues. ◆ Verify port speeds and switch port speeds for the MirrorView ports. ◆ Determine if the connectivity loss is on one or both SPs. ◆ Verify that the storage system or SP is operating properly with no other faults or symptoms. ◆ Try to disable and re-enable the MirrorView connection from both storage systems.

Table 4. Summary of MirrorView/A failures *(continued)*

Event	Result and recovery
Failures when adding secondary images	Make sure that: <ul style="list-style-type: none"> ◆ The connection between storage systems works. ◆ You are managing both storage systems (which may require managing two domains). ◆ The secondary LUN is available and the same size as the primary image. ◆ The secondary image LUN is not in the storage group. ◆ The secondary image LUN is not already a secondary image, of either a synchronous or asynchronous mirror. ◆ The reserved LUN pools on both primary and secondary storage systems are adequately configured. ◆ The secondary LUN is not set up as a destination for SAN Copy.
When the secondary image cannot be synchronized	Make sure that: <ul style="list-style-type: none"> ◆ The connection between the storage systems is still good. ◆ The recovery policy is set to auto, not manual. ◆ The secondary SP is working. ◆ If the image is fractured, view the Last Image Error. Try manually fracturing and then manually synchronizing the image.

Recovering from serious errors

In the unlikely event that the mechanism for tracking changes made to the primary image fails (for example, insufficient memory is available on the SP), the secondary image is marked as permanently fractured. To recover from this situation, you must remove the secondary image from the mirror, and then add it again (which does a full resynchronization). This

failure may indicate that you are using close to the storage system’s capacity for layered features.

Some other serious failures will transition MirrorView/A into a degraded mode of operation, where administrative requests will be rejected and no further updates run. Degraded mode affects only a single SP; the other SP in the storage system may continue to run normally (depending on the nature of the failure).

When an SP enters degraded mode, the system logs an event that indicates why MirrorView/A is in the degraded mode. Usually you can recover from the degraded mode by simply rebooting the affected SP, but some specific cases require you to check other components that MirrorView/A uses before rebooting the SP. [Table 5 on page 104](#) lists various scenarios in which MirrorView/A goes to the degraded mode and the recovery options you can take.

Table 5. Recovery from degraded mode

Event	Result and recovery
Internal memory corruption	Mirror data does not match the expected value. Reboot the SP.
Serious, unexpected errors	MirrorView/A receives unexpected errors from its underlying components during operation. Look at the event log for a record of errors, and take steps to correct them. For example, if the reserved LUN pool LUNs are faulted, recover the reserved LUN pool LUNs. Then reboot the SP.
Internal fracture failure	A fracture operation fails due to reasons other than an error you made. Look at the event log for the appropriate failure reason. Reboot the SP to fix the problem.

How consistency groups handle failures

When a failure occurs during normal operations for consistency groups, MirrorView/A lets you perform several actions in order to recover.

When recovering from failures, MirrorView/A achieves three goals:

- ◆ Preserves data integrity
- ◆ Minimizes the amount of time that data is unavailable to the user
- ◆ Ensures that the consistency of the consistency group is maintained

Access to the SP fails

Consider a consistency group that has member mirrors, some of which SP A controls and some of which SP B controls. If SP A on the primary storage system fails, then software on the attached server, for example, PowerPath, moves control of the mirrors that were controlled by SP A to SP B. This allows applications on the server, as well as the mirroring of data to the secondary storage system, to continue uninterrupted. However, as part of the transfer of control, the consistency group becomes system fractured. If the recovery policy is set to automatic, an update automatically starts on the surviving SP (SP B in this example). However, if the recovery policy is manual, you must manually start an update.

Primary storage system fails

If the storage system running the primary consistency group fails, access to the data in the group's member LUNs is lost. You can either repair the failed storage system and then continue operations, or you can promote the secondary consistency group, and access the data from the secondary storage system.

iSCSI-related failure conditions

Table 6 on page 106 shows the iSCSI-related failure conditions and the reasons they occur.

Table 6. iSCSI-related failures

Failure condition	Reason
You selected iSCSI and Fibre Channel connections for the same storage system and clicked Enable in the Manage Mirror Connections dialog box.	You selected iSCSI and fibre connections for the storage system; you must choose one or the other.
You selected iSCSI for a storage system already configured to be reached with a Fibre Channel for MirrorView connection.	You tried to connect storage system A to storage system B with an iSCSI connection, but a Fibre Channel connection between these two storage systems already exists. If you want to change from Fibre Channel to iSCSI, you must first disable the Fibre Channel connection to storage system B.
You selected Fibre Channel for a storage system already configured to be reached with an iSCSI for MirrorView connection.	You tried to connect storage system A to storage system B with a Fibre Channel connection, but an iSCSI connection between these two storage systems already exists. If you want to change from iSCSI to Fibre Channel, you must first disable the iSCSI connection to storage system B.

Appendix B

Error Codes

This appendix lists the MirrorView/A error codes returned by the CLI. The numbers proceed sequentially.

- ◆ [Error codes on page 108](#)

Error codes

Error value	Description
1	Invalid command line parameters
2	Too few parameters.
3	Too many parameters.
5	Invalid LUN number.
30	LUN does not exist.
39	Cannot access device.
41	Command not supported.
66	Error returned from agent.
89	ToInvalid switch combination.
123	Snapshot does not exist.
125	Invalid SP name.
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 r1 (RAID1), r3 (RAID 3), r5 (RAID 5), or r1_0 (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 initializearray 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.
145	Mirror not found.

Error value	Description
146	Image not found.
147	Synchronized rate cannot be changed since input image is primary.
148	Cannot specify both -uid and -filename .
149	Invalid cache.
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 -filename.
156	Path to store files is required.
157	Cannot specify both -all and -filename.
158	Enter file index or <code>quit</code> .
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.
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.

Error Codes

Error value	Description
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	Email page response test failed.
187	Modem page response test failed.
188	Snmp response test failed.
189	Phone home response test failed.
190	Mandatory switch for email/emailpage.
191	Mandatory switch for modem page.
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 be opened.
196	Specified user already exists.
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 server.
208	Switch -pathname must be specified.
209	Get local server attributes failed.
210	This version of FLARE Software does not support Hi5 RAID type.

Error value	Description
211	Only one of the switches -snapshotid, -snapshotname, or -lun can be specified.
212	Specified session and snapshot must be based on the same source LUN.
215	The HLU (server 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.
223	Allocate log needs to specify a LUN with a valid owner.
224	This request has been issued through the SP that is not the current owner of the targeted LUN.
225	Simulated/non-simulated sessions are mixed on the same targeted LUN.
234	The LUN is in a storage group.
262	This version of FLARE software does not support rollback.
263	The switch -rate must be specified.
264	Invalid rollback rate. Valid values are high, medium and low.
265	Cannot change the rollback rate. The specified session is not rolling back.
266	Cannot start a rollback operation. The specified session is not persistent.
267	Cannot start a rollback operation. The specified session is rolling back
268	Cannot start a rollback operation. At least one source LUN is involved in another rollback operation.
269	Cannot create a recovery session. At least one source LUN has maximum sessions already.
270	Invalid SnapView session name.
271	Simulation mode is not supported in this revision.
272	Snapshot cache chunk size cannot be changed in this revision.
273	Cannot stop a session when it is rolling back.
274	Cannot activate a snapshot on a session when the session is rolling back.
275	Cannot deactivate a snapshot when it is activated on a session that is rolling back.
276	Multiple mirrors with the same name have been found. Please use UID for a specific mirror.
294	This request must be issued from the SP that owns the session.
295	This request must be issued from the SP that owns the session that the snapshot is activated on.
296	Snapshot is already inactive.

Error Codes

Error value	Description
297	Snapshot name with leading or trailing white space is invalid.
298	Cannot create a snapshot on private LUNs (hot spare, snap cache, PSM, and so on).
337	Remote mirror write intent log size is too small.
352	Consistent mode is not supported in this revision.
414	Mirror cannot resolve path.
433	LUN is in use.
0x6201	Poll failed
0x6202	Cannot create a session with a name longer than 255 characters.
0x6203	Cannot start a session on a LUN that has another session Rolling Back.
0x6205	The Snapshot name is too long
0x6207	Could not create the Snapshot because the SnapCopy driver could not be added to the Source LUN's driver stack.
0x6209	The attribute(s) cannot be set for private lun(s).
0x620A	Could not add SnapView driver to the LUN stack.
0x620B	Illegal default owner specified.
0x620C	This SnapCopy session is not currently Rolling Back.
0x620D	A LUN for this SnapCopy session is currently being Rolled Back.
0x620E	SnapCopy RollBack is not supported.
0x620F	A Mirror or Clone is preventing a RollBack of this session. You must administratively fracture a Mirror or Clone when it is consistent or synchronized before starting a RollBack.
0x6210	Snapshot is not associated with a source LUN.
0x6211	Session not found in session list.
0x6212	Snap Copy Session is transitioning.
0x6213	Snap LUN region not found.
0x6214	Too many sessions for this snapshot.
0x6216	Could not remove the SAN Copy driver from this Snapshot.
0x6217	The LU is in a storage group.
0x6218	Snapshot is already inactive.
0x6219	No LUs to add to Reserved LUN pool (OR) Trying to exceed the maximum number of LUNs supported (OR) LU already in Reserved LUN pool (OR) The LU is transitioning (OR) LU is in a storage group (OR) LU in use by array feature.
0x621A	LU cannot be removed because it is currently allocated. (OR) LU not in Reserved LUN pool (OR)

Error value	Description
0x621F	There is an active SAN Copy session operating on a LUN in the SnapView session. Please wait for it to complete before attempting to RollBack this session.
0x6221	LU is not available.

Error codes - Hexadecimal

Error value (Decimal)	Error value (Hexadecimal)	Description
19873	0x4DA1	Invalid Async Mirror command.
19875	0x4DA3	Async feature object is not available.
19877	0x4DA5	Async enabling software is not installed.
19879	0x4DA7	Unexpected Async error occurred.
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 errors occurred.
19886	0x4DAE	Intentlog allocate error.
19887	0x4DAF	Intentlog deallocate error.
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 cannot be chosen as mirror LU.

Error value (Decimal)	Error value (Hexadecimal)	Description
19900	0x4DBC	Cannot perform create, add secondary image or allocate WIL on a LU while it is migrating.
19901	0x4DBD	Cannot perform setfeature on a LU while it is private.
19902	0x4DBE	Moved from CLIErr.h.
19903	0x4DBF	setfeature command is not supported on a thin LUN.
Consistency group error codes		
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.
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.

Secure CLI Command Coverage

This appendix provides a list of commands that are supported by Secure CLI on various storage systems. Unless otherwise noted once supported all commands will be supported on all subsequent releases. The commands that are not supported on specific systems are also listed.

- ◆ [Secure CLI command coverage on page 116](#)

Secure CLI command coverage

Commands	CX series	CX3 series	CX4 series
<code>mirror -async -addimage</code>	6.19	6.22	6.28
<code>mirror -async -change</code>	6.19	6.22	6.28
<code>mirror -async -changeimage</code>	6.19	6.22	6.28
<code>mirror -async -create</code>	6.19	6.22	6.28
<code>mirror -async -destroy</code>	6.19	6.22	6.28
<code>mirror -async -fractureimage</code>	6.19	6.22	6.28
<code>mirror -async -info</code>	6.19	6.22	6.28
<code>mirror -async -list</code>	6.19	6.22	6.28
<code>mirror -async -promoteimage</code>	6.19	6.22	6.28
<code>mirror -async -removeimage</code>	6.19	6.22	6.28
<code>mirror -async -setfeature</code>	6.19	6.22	6.28
<code>mirror -async -syncimage</code>	6.19	6.22	6.28
<code>mirror -disablepath</code>	6.19	6.22	6.28
<code>mirror -enablepath</code>	6.19	6.22	6.28
<code>mirror -async -creategroup</code>	6.19	6.22	6.28
<code>mirror -async -destroygroup</code>	6.19	6.22	6.28
<code>mirror -async -addtogroup</code>	6.19	6.22	6.28
<code>mirror -async -removefromgroup</code>	6.19	6.22	6.28
<code>mirror -async -changegroup</code>	6.19	6.22	6.28
<code>mirror -async -syncgroup</code>	6.19	6.22	6.28
<code>mirror -async -fracturegroup</code>	6.19	6.22	6.28
<code>mirror -async -promotegroup</code>	6.19	6.22	6.28
<code>mirror -async -listgroups</code>	6.19	6.22	6.28

Note: The version number indicates the earliest Navisphere version that supported the command for each system type.

A

active state

State in which a remote mirror is running normally.

attention state

State in which the mirror's secondary image is fractured and the mirror is configured to generate an alert. The mirror continues to accept server I/O in this state.

auto recovery

MirrorView option to have synchronization start as soon as a system-fractured secondary image can be reached.

C

consistency group

A set of asynchronous mirrors that are managed as a single entity and whose secondary images always remain in a consistent and recoverable state with respect to their primary image and each other.

consistency group condition

Detailed information about the consistency group, including whether the group is normal, initializing, updating, admin fractured, system fractured, or waiting on admin.

consistency group state

The current state of the consistency group: Synchronized, Consistent, Quasi-Consistent, Synchronizing, Out-of-Sync, Scrambled, Rolling Back, or Empty.

consistent state (of image)

State in which a secondary image is identical to either the current primary image or to some previous (consistent point-in-time) instance of the primary image.

F**fracture**

A condition in which I/O does not propagate to the secondary image; it can result when you initiate the fracture (*admin fracture*) or when the system determines that the secondary image is unreachable (*system fracture*). An admin fracture may also occur if the MirrorView/A software detects an error condition that requires administrative intervention to correct.

I**image condition**

The condition of a secondary image that provides additional information about the status of updates for the image. Values include normal, administratively fractured, system fractured, updating, synchronizing, or waiting on admin.

image state

The indication of the relationship between a secondary image and the primary image of a mirror. The image states are: Synchronized, Consistent, Synchronizing, Out-of-Sync, and Rolling Back.

O**out-of-sync state**

A remote mirror state in which the software does not know how the primary and secondary images differ; therefore a full synchronization is required to make the secondary image(s) usable for recovery.

Also see *image state*.

P**primary image**

The LUN on the production storage system that contains user data and is the source for data copied to the secondary image. There is one primary image and zero or one secondary image. A remote mirror is ineffective for recovery unless it has a secondary image. This manual also refers to primary image as primary.

promote (to primary)

The operation by which the administrator changes an image's role from secondary to primary. As part of this operation, the previous primary image becomes a secondary image. If the previous primary image is unavailable when you promote the secondary image (perhaps because the primary site suffered a disaster), the software does not include it as a secondary image in the new mirror. A secondary image can be promoted if it is in either the Synchronized state or the Consistent state and an update is not currently transferring data.

R**recovery time objective**

A time interval you set for MirrorView/A to automatically update the secondary image. You can specify it as the time interval between the start of two consecutive updates or between the end of the last update and the beginning of the next update.

remote mirror

The combination of a LUN on one storage system, called the primary image, and another LUN on a different storage system, called the secondary image. The software maintains the secondary image as an exact copy of the primary image at some (possibly previous) point in time. If the server and/or storage system at the primary site fails, you can promote the secondary image to take over the role of the primary, thus allowing continued access to your production data.

remote mirror state

Active and attention. In the active state, the remote mirror is running normally. In the attention state, a site or communications failure has caused several accessible images to fall below the required minimum. Administrator action is required to correct the issue with a secondary image.

rolling back state

The state that occurs after a successful promotion where there was an unfinished update to the secondary image. This state persists until the rollback operation completes.

S**secondary image**

A LUN that contains a mirror of the primary image LUN. There can be zero or one secondary image. This manual also refers to secondary image as secondary.

states

Remote mirror states and image states. The remote mirror states are: Active and Attention. The image states are: Synchronized, Consistent, Synchronizing, Out-of-Sync, and Rolling Back.

synchronizing state

A secondary image in the process of initial synchronization. The data in the secondary image is not usable for recovery until the synchronization operation completes. Thus, you cannot promote an image in the Synchronizing state to the primary image. Once initial synchronization completes, a MirrorView/A mirror should not return to this state, but should stay consistent or synchronized.

U**unfinished update**

State in which the secondary image is fractured when the data is being transferred as well as not being transferred.

update in progress

A process where data is currently being transferred to the secondary storage system.

A

account, user 30
 application integration 23

B

bidirectional mirroring 21

C

case (of characters in CLI commands) 30
 CLI
 MirrorView/A subcommands 25
 Consistency Group
 changing properties 83
 creating 74
 fracturing images in 87
 updating 85

E

error
 codes for MirrorView/A subcommands 107
 error codes for SnapView subcommands 108

F

failures, how MirrorView handles 96
 file security 30

I

integration with EMC SAN Copy software 22
 integration with EMC SnapView Snapshot software 22

integration with EMC SnapView software 22

M

mirror changegroup subcommand 83
 mirror creategroup subcommand 74
 mirror enablepath subcommand 54
 mirror fracturegroup subcommand 87
 mirror fractureimage subcommand 56
 mirror info subcommand 58
 mirror list subcommand 60
 mirror promoteimage subcommand 63
 mirror removeimage subcommand 66
 mirror syncgroup subcommand 85
 MirrorView
 handling failures 96
 MirrorView/A
 about Chapter 2 25
 configuration guidelines 15
 operations overview 27
 overview 14
 prerequisites 14
 MirrorView/A error codes 107
 MirrorView/A subcommands 25
 MirrorView/A subcommands, Chapter 2 25

P

password 30
 path
 remote mirror, enabling 55
 primary image 14, 96
 failure 96
 promoting
 a secondary image to a primary 96
 promotion

promotion (*continued*)
of secondary image 17

R

remote mirror
fracturing image 56
info command 58
listing mirror information 60
path, enabling 54
promoting image 63
removing image 66
remote mirroring, promotion 17
replication over long distances 23
requirements, Secure CLI 30

S

secondary image
failure 98
promotion 17
Secure CLI 30

Secure CLI (*continued*)
requirements 30
using security 30
session, MirrorView/A, see MirrorView/A 27
SnapView
error codes 108
subcommands
mirror changegroup 83
mirror creategroup 74
mirror enablepath 54
mirror fracturegroup 87
mirror fractureimage 56
mirror info 58
mirror list 60
mirror promoteimage 63
mirror removeimage 66
mirror syncgroup 85

U

Unisphere security 30
user security 30
username 30