

**EMC Navisphere
Manager**

ADMINISTRATOR'S GUIDE

**P/N 069001125
REV A12**

EMC Corporation

Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.EMC.com

Copyright © 2003-2005 EMC Corporation. All rights reserved.

Published February, 2006

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.

Trademark Information

EMC², EMC, CLARiiON, Navisphere, PowerPath, Symmetrix, and TimeFinder are registered trademarks, and Access Logix, FLARE, Powerlink, SAN Copy, MirrorView, and SnapView are trademarks of EMC Corporation.

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

Preface	xvii
----------------------	------

Chapter 1 About EMC Navisphere Manager

Terminology.....	1-2
About EMC Navisphere software	1-5
Navisphere Storage-System Initialization Utility	1-5
Navisphere Host and SP Agents	1-6
Navisphere Server Utility	1-6
Navisphere Manager.....	1-8
Navisphere Integrator.....	1-8
Navisphere Storage Management Server.....	1-9
Navisphere CLI	1-9
Navisphere Analyzer	1-11
Navisphere environments	1-11
Storage-system configuration and management.....	1-12
To set up Navisphere security and domains	1-12
To install optional software on a CX-series or FC-series storage system	1-12
To configure a shared storage system with Manager	1-12
To manage a shared storage system with Manager	1-13

Chapter 2 Preparing to Use and Starting Manager

Preparing to use Manager	2-2
Installing a new version of the JRE	2-3
Starting a Manager session.....	2-3
Selecting a new language for Navisphere Manager.....	2-5

Chapter 3 Installing Storage-System Software and Disk Firmware

Storage-system software installation introduction	3-2
Software installation wizard overview	3-3
Starting the software installation wizard	3-4
Entering non-disruptive upgrade delay	3-4
Selecting software	3-5
Pre-install validation	3-6
Viewing installation summary	3-9
Viewing and monitoring installation status	3-9
Committing (finalizing) the software installation	3-11
Reverting back to the previous software version	3-12
Displaying software status for CX-series and FC4700-series storage systems	3-12
Using the newly installed software packages	3-13
Installing and upgrading software on pre-FC4700 storage systems	3-14

Chapter 4 Managing and Monitoring AX-Series Storage Systems

Managing AX-series systems with Navisphere Manager	4-2
Managing Prerequisites	4-2
Navisphere Manager behavior with AX-series systems	4-4
AX-series to CX-series terminology differences	4-4
Status reporting differences	4-5
AX-series tree structure	4-6
Right-click menu options for AX-series storage systems ...	4-7
Dialog box limitations	4-11
Managing and monitoring configurations	4-13
Managing and monitoring in a single-domain environment	4-13
Managing and monitoring in a multi-domain environment	4-13
Managing and monitoring in a multi-system domain	4-14

Chapter 5 Setting Storage System Failover Properties

Failover software values overview	5-2
Setting failover values using the Failover Setup Wizard	5-3
Starting and using the Failover Setup Wizard	5-4
Setting failover values for initiators connected to a specific storage system	5-10
Setting failover values when registering a new or existing initiator	5-10

Chapter 6 **Configuring the Remote Host Agent**

Configuring Host Agents remotely (pre-FC4700 series)	6-2
Adding privileged users	6-3
Scanning for devices	6-3
Updating the communications channels list	6-5
Updating parameters	6-6

Chapter 7 **Setting Storage-System Properties**

Setting storage-system management control properties	7-2
Management controls	7-2
Storage-system management control properties	7-3
Enabling management control for a pre-FC4700 shared storage system	7-4
Enabling and disabling management control for servers ..	7-5
Setting storage-system general configuration properties	7-7
General configuration properties	7-7
Setting storage-system memory properties	7-9
Assigning memory to partitions	7-10
Setting storage-system cache properties	7-11
Storage-system cache properties	7-12
Setting the cache properties	7-15
Setting the fair access property (pre-FC4700 only)	7-16
Fair access to the storage-system resources	7-16
Enabling fair access to a storage system	7-16
Setting the SP network and ALPA properties (CX-series or FC4700-series only)	7-17
Setting the SP network properties	7-17
Setting the SP ALPA (Arbitrated Loop Physical Address) properties	7-18
Setting the SP fibre speed property	7-19
Setting the SP port speed	7-19
Adding privileged users to the SP agent configuration file...	7-20
Setting the battery test time	7-20
Registering HBAs and setting HBA initiator record values ...	7-21
Registering HBAs with a specific storage system	7-21
Registering an existing unregistered HBA	7-22
Registering a new unregistered HBA	7-23
Registering HBAs from a specific host	7-23
To deregister an HBA initiator	7-25

Chapter 8 **Creating LUNs and RAID Groups**

About LUNs	8-2
LUNs	8-2
RAID types	8-2
Number of disks in a LUN.....	8-4
Disks that cannot be hot spares	8-4
LUN properties.....	8-5
About RAID groups	8-11
RAID group expansion.....	8-11
RAID group defragmentation	8-11
Disk replacement	8-11
RAID group properties.....	8-12
RAID group characteristics	8-13
Creating RAID groups	8-14
To create a RAID group	8-14
Creating LUNs on RAID groups	8-16
To create LUNs on a RAID group	8-17
Binding a new LUN with the exact size of an existing LUN...	8-20
Creating LUNs on a non-RAID group storage system	8-21
To bind LUNs on a non-RAID group storage system.....	8-21

Chapter 9 **Setting Up Access Logix and Creating Storage Groups**

Setting the storage-system data access property.....	9-2
Data access control and storage groups	9-2
Enabling data access control for a storage system.....	9-2
Creating storage groups.....	9-3
Adding LUNs to a storage group	9-4
Connecting a server to a storage group.....	9-5
Connecting a server to a storage group	9-6
Verifying server connections to a storage group.....	9-6
Verifying or editing device information in the host agent configuration file (pre-FC4700 storage systems).....	9-9
Verifying or editing host agent device information on an AIX, HP-UX, Linux, or NetWare or Windows server	9-9
Verifying or editing agent device information on a Solaris server	9-10

Chapter 10 Expanding LUN Capacity With MetaLUNs

MetaLUN terminology	10-2
MetaLUNs overview	10-3
Expanding a LUN	10-5
Restrictions	10-5
Results	10-6
Expanding a metaLUN	10-9
Restrictions	10-9
Results	10-10
Using the expand storage wizard to expand LUN capacity ..	10-12
Starting the Expand Storage Wizard	10-12
Selecting storage expansion type	10-13
Selecting expansion LUNs	10-13
Specifying metaLUN capacity	10-14
Specifying metaLUN settings	10-15
Viewing metaLUN information	10-17
Displaying and changing MetaLUN properties	10-17
Monitoring expansion status	10-19
Destroying a metaLUN	10-19
Using the added capacity of an expanded SnapView, MirrorView / A, MirrorView / S or SAN Copy LUN	10-20

Chapter 11 Managing the Reserved LUN Pool for LUN Replication

Introduction to the reserved LUN pool	11-2
Allocation of reserved LUNs	11-2
How SPs manage the reserved LUN pool	11-3
Estimating the reserved LUN pool size	11-3
Estimating a suitable reserved LUN pool size	11-4
Configuring the reserved LUN pool	11-5
Monitoring reserved LUN pool usage	11-7
Displaying reserved LUN pool properties	11-8

Chapter 12 Setting Up and Using Event Monitor

Event Monitor overview	12-2
Event Monitor terminology	12-4
Distributed and centralized monitoring	12-5
Monitoring guidelines	12-5
Configuring distributed monitoring	12-5
Configuring centralized monitoring for a single domain	12-7
Configuring centralized monitoring for multiple domains	12-9

Event Monitor configuration overview	12-10
Configuring Event Monitor using the wizard	12-11
Configuring Event Monitor using menu options	12-17
Adding monitored Agents	12-17
Removing monitored Agents	12-18
Configuring log file size and network timeout	12-18
Creating new templates without the wizard	12-19
Setting up responses for the templates	12-21
Creating templates with customized responses	12-25
Creating new templates based on existing templates	12-26
Testing responses	12-27
Testing configurations	12-28
Using insert test event	12-28
Viewing test events	12-28
Viewing responses to test events	12-29
Deleting customized responses	12-29
Importing and exporting templates	12-29
Importing templates	12-29
Exporting templates	12-30
Viewing and modifying template properties	12-31
Viewing template properties	12-31
Modifying template properties	12-31
Adding new threshold properties to converted templates	12-32
Deleting templates	12-32
Applying templates to storage systems	12-33
Using Event Monitor	12-34
Understanding SP event codes and event severity	12-34
CX-Series and FC4700-Series code format	12-36
Viewing system event logs	12-36
Viewing details about a specific event	12-37
Viewing the response log	12-37
Viewing the message file	12-38
Saving events to the log file	12-38
Printing events	12-38
Clearing events in the events window	12-39

Chapter 13 **Monitoring Storage-System Operation**

Monitoring storage-system operation.....	13-2
Manual polling	13-3
Storage-system faults.....	13-4
Faulted disk icon	13-5
Faulted SP icon	13-8
Faulted LCC icon	13-9
Faulted fan A or fan B icon	13-10
Faulted power supply or VSC icon	13-12
Faulted SPS icon	13-13
Faulted BBU icon.....	13-14
Displaying storage-system component and server status.....	13-15
Using the SP event viewer for monitoring SP events	13-17
Displaying the event logs for an FC4700 SP	13-17
Filtering events	13-18
Viewing event details	13-18
Saving events to a log file	13-19
Printing events.....	13-19
Clearing events in the events window.....	13-19
Generating and transferring SP log files.....	13-20
Monitoring storage-system connectivity status.....	13-21
Updating host connectivity status	13-23
Monitoring the status of trespassed LUNs.....	13-24
Monitoring domain status	13-24
Displaying NAS device status.....	13-25
To display NAS device properties.....	13-25

Chapter 14 **Reconfiguring LUNs, RAID Groups, and Storage Groups**

Reconfiguring LUNs.....	14-2
To change a LUN's unique ID or element size	14-2
To change the RAID type of a LUN	14-3
To change the user capacity of a LUN	14-3
Changing the remaining properties of a LUN.....	14-4
Changing LUN prefetch (read caching) properties	14-5
Transferring the default ownership of a LUN.....	14-8
Unbinding a LUN	14-10
Using the LUN migration feature to copy a LUN.....	14-13
LUN migration requirements.....	14-14
LUN migration limits	14-14
Access to migration objects	14-14
Migration states	14-15

Migrating a LUN	14-15
Migration status - LUN properties.....	14-16
Canceling a migration.....	14-16
Displaying a migration summary	14-16
Reconfiguring RAID groups	14-17
Changing the expansion/defragmentation priority or the automatically destroy after last LUN unbound property of a RAID group	14-17
Expanding a RAID group.....	14-18
Defragmenting a RAID group	14-20
Destroying a RAID group	14-21
Reconfiguring storage groups.....	14-23
Changing the name of a storage group	14-23
Adding or removing LUNs from storage groups.....	14-23
Connecting servers to a storage group or disconnecting servers from a storage group	14-24
Destroying storage groups.....	14-25

Chapter 15 Reconfiguring Storage Systems

Upgrading a storage system to support caching	15-2
Installing the hardware components for caching	15-2
Setting up caching	15-3
Replacing disks with higher capacity disks.....	15-4
Connecting a new server to a shared storage system.....	15-6
Disconnecting a server from a shared storage system	15-6

Chapter 16 Managing iSCSI (Internet SCSI) Storage Systems

About iSCSI storage systems in IP networks.....	16-2
Using iSNS (Internet Storage Naming Service) within an iSCSI storage environment	16-5
Setting up an iSCSI storage system and server for I/O	16-7
Configuring storage system iSCSI port network parameters	16-9
Configuring the iSNS storage environment	16-10
Configuring iSCSI server initiator network parameters.....	16-12
Testing iSCSI network connections	16-13
Registering server HBAs or NICs	16-13
Setting up CHAP credentials on the storage system.....	16-14
Enabling CHAP authentication on the storage system.....	16-16
Managing CHAP credentials	16-17
Modifying CHAP credentials	16-18

Chapter 17 Managing Legacy Storage Systems and NAS Systems

Portal management architecture	17-2
Adding and configuring a portal.....	17-5
Adding a portal	17-5
Assigning storage systems and NAS systems to portals ..	17-6
Monitoring portal system status.....	17-8
Adding IP addresses to the agent configuration files.....	17-8
Adding IP addresses to the Host Agent configuration file.....	17-9
Adding IP addresses to the SP Agent configuration file	17-11

Appendix A Using Trees and the Main Window

Introduction to tree views	A-2
Specifying the managed storage systems to display	A-6
Selecting tree icons.....	A-7
Icon indicators	A-7
Right-click menu options for basic tree icons	A-8
Main window	A-22
Application icon	A-23
Menu bar	A-23
Toolbar	A-25
Status bar.....	A-26
Workspace	A-26

Index	i-1
-------------	-----

3-1	Enablers for AX-series and CX-series storage systems	3-3
3-2	Validation checks	3-6
3-3	Validation check icons	3-8
3-4	Database disks for different storage-system types	3-14
7-1	Hardware requirements for write caching	7-11
8-1	Number of disks you can use in RAID types	8-4
8-2	Disks that cannot be hot spares	8-4
8-3	Valid rebuild priorities	8-6
8-4	LUN properties available for different RAID types	8-9
8-5	Default LUN property Values for different RAID types	8-10
8-6	Recommendations for creating LUNs with different RAID types	8-10
8-7	Default RAID group property values	8-12
8-8	Maximum number of LUNs per RAID group	8-12
10-1	The effects of expanding a LUN	10-7
10-2	The effects of expanding a metaLUN	10-11
12-1	Centralized monitoring response matrix	12-3
12-2	Distributed monitoring response matrix - host resident	12-3
12-3	Distributed monitoring response matrix - SP resident	12-3
13-1	Application icon states	13-2
13-2	Disk failure states	13-5
13-3	Cache vault disks	13-8
13-4	Number of fans in a C-series storage system	13-11
13-5	BBU failure states	13-14
14-1	Default prefetch properties values	14-6
14-2	Available prefetch properties values - general	14-7
14-3	Available prefetch properties values - constant	14-7
14-4	Available prefetch properties values - variable	14-7
14-5	LUN migration limits	14-14
14-6	Number of disks you can use in RAID types	14-18
15-1	Hardware requirements for write caching	15-2

15-2	Database disks	15-4
16-1	iSCSI CHAP management window	16-17
A-1	Filters for displaying managed storage systems	A-6
A-2	Storage-system status icons	A-7
A-3	Basic storage tree icons: images and descriptions	A-9
A-4	Basic hosts tree icons: images and descriptions	A-19
A-5	Basic monitors tree icons: images and descriptions	A-21
A-6	Toolbar icon image	A-25

10-1	Sample stripe expansion	10-4
10-2	Sample concatenate expansion	10-4
12-1	Distributed monitoring system	12-6
12-2	Centralized monitoring system	12-8
16-1	Nodes - initiator and target	16-2
16-2	FC and iSCSI storage systems in an IP network	16-4
16-3	Sample iSNS storage configuration	16-6
17-1	Sample storage system portal environment	17-3
17-2	Sample server portal environment	17-4
A-1	Sample partially expanded storage tree	A-3
A-2	Sample partially expanded hosts tree	A-4
A-3	Sample partially expanded monitors tree	A-5
A-4	Main window	A-22

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

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

About this manual

This guide describes the tasks for setting up, configuring, and managing a storage system using EMC Navisphere® Manager 6.X software. Each major section includes introductory information and a general procedure for completing a task. This manual is not intended for use during the actual setup, configuration, and management of storage systems so the steps in the procedures purposely do not include screen captures of the dialog boxes.

The introductory information and *detailed* steps for each procedure appear in the Navisphere Manager online help so you have the complete information available when you actually set up, configure, and manage storage systems, should you require help.

Audience

This guide is part of the EMC Navisphere Manager documentation set, and is intended for use by customers and service providers who use EMC Navisphere Manager 6.X to configure and manage storage systems.

This guide applies to Manager Base and Manager.

Organization

This guide is organized as follows:

- | | |
|------------|--|
| Chapter 1 | Introduces EMC Navisphere Manager 6.X, and outlines the tasks you will perform to configure and manage storage systems. |
| Chapter 2 | Explains how to start an EMC Navisphere Manager 6.X session. |
| Chapter 3 | Describes how to use the wizard to install software bundles on CX-Series and FC4700-Series storage systems, and how to install and update Core Software on non-FC4700 storage systems. |
| Chapter 4 | Describes the requirements for managing and monitoring AX-Series storage systems using Navisphere Manager. |
| Chapter 5 | Introduces storage system failover values and explains how and when to set them. |
| Chapter 6 | Explains how to configure the Remote Host Agents for pre-FC4700 storage systems. |
| Chapter 7 | Explains how to set the general, memory, and cache properties; how to set the SP and ALPA properties; and how to set the battery test time for a storage system. |
| Chapter 8 | Explains how to create logical units (LUNs) that are composed of disks. This chapter also explains how to create RAID Groups (and the LUNs in the RAID groups). |
| Chapter 9 | Explains how to enable and use the Access Logix™ feature — software that lets you create Storage Groups on a storage system so multiple hosts can have their own LUNs. |
| Chapter 10 | Explains how to expand LUN capacity with metaLUNs. |
| Chapter 11 | Explains how to configure the reserved LUN pool which is required for running replication software such as SnapView MirrorView / A, and Incremental SAN Copy. |

Chapter 12	Introduces the Event Monitor feature, which provides the ability to track the status of storage systems connected to servers on a direct network or storage area network (SAN). This chapter explains how to configure and use the Event Monitor feature.
Chapter 13	Describes how EMC Navisphere Manager 6.X displays status information. This chapter also explains how to monitor storage-system operation and display storage-system event messages.
Chapter 14	Explains how to reconfigure LUNs, RAID Groups, and Storage Groups.
Chapter 15	Explains how to reconfigure storage systems.
Chapter 16	Describes how to configure iSCSI (IP SCSI) storage systems in an IP storage network.
Chapter 17	Explains how to assign and configure portal systems so that you can manage and configure legacy storage systems and NAS systems.
Appendix A	Describes the EMC Navisphere storage-system trees and the Main window that you use for monitoring and configuring storage systems.

Related documents

Other related documents include:

- ◆ *EMC Navisphere Security Administrator's Guide* (P/N 069001124)
- ◆ *EMC MirrorView for EMC Navisphere Administrator's Guide* (P/N 069001161)
- ◆ *EMC SnapView for EMC Navisphere Administrator's Guide* (P/N 069001180)
- ◆ *EMC SAN Copy for EMC Navisphere Administrator's Guide* (P/N 069001188)
- ◆ *EMC Navisphere Command Line Interface Reference* (P/N 069001038)
- ◆ *EMC Navisphere Manager Release Notes* (P/N 085090645)

Conventions used in this manual

EMC uses the following conventions for notes and cautions.

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



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Typographical conventions

This guide uses the following format conventions

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

Finding current information

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

To access EMC Powerlink, use the following link:

<http://powerlink.emc.com>

After you log in, select **Support > Document Library** and find the following:

- ◆ *EMC Navisphere Manager Release Notes*, P/N 085090645
- ◆ The latest version of this guide that is applicable to your software revision.
- ◆ *EMC Installation Roadmap for CX-Series, AX-Series, and FC-Series Storage Systems*, which provides a checklist of the tasks that you must complete to install your storage system in a storage area network (SAN) or direct attach configuration.

Using online help

The following help is available online from the Navisphere Manager user interface.:

- ◆ A set of organized, linked help topics
To access the online help table of contents, click **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.

Where to get help

For questions about technical support and service, contact your service provider.

If you have a valid EMC service contract, contact EMC Customer Service at:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

Follow the voice menu prompts to open a service call, then select CLARiiON Product Support.

Sales and customer service contacts

For the list of EMC sales locations, please access the EMC home page at:

<http://www.EMC.com/contact/>

For additional information on the EMC products and services available to customers and partners, refer to the EMC Powerlink Web site at:

<http://Powerlink.EMC.com>

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to techpub_comments@EMC.com with your opinions of this guide.

About EMC Navisphere Manager

EMC Navisphere® Manager provides a web-based user interface that lets you configure and manage disk-array storage systems connected to hosts on a network.

This chapter briefly describes the following:

- ◆ Terminology 1-2
- ◆ About EMC Navisphere software 1-5
- ◆ Storage-system configuration and management..... 1-12

Terminology

Term	Meaning
Access Logix™ option	See shared storage system .
Analyzer	EMC Navisphere Analyzer (a performance analyzer).
AX-Series storage systems	AX100SC, AX100SCi, AX100 and AX100i storage systems where AX100SC and AX100SCi are single SP storage systems; AX100 and AX100i are dual SP storage systems; AX100 and AX100SC are Fibre Channel storage systems; AX100SCi and AX100i are iSCSI storage systems.
C-series storage system	A C3000, C2x00, C1900, or C1000 series storage system.
CLI	EMC Navisphere Command Line Interface.
CX-Series storage system	CX700, CX600, CX500, CX400, CX300, CX200, and CX200LC
client	A host (computer or laptop) that has an Internet browser and connects to a storage application server via network. You use it to manage connected storage systems.
directory	In the context of the web-based Manager, a directory is the database that contains information such as SP IP addresses of the storage systems and the accounts of global users in the domain. The software creates this directory when you create a domain.
domain	In the context of the web-based management software, a domain is a group of storage systems that you can manage from a single management application session. You can divide your storage systems into multiple domains as long as each storage system is in only one domain. Each domain has a directory that defines the storage systems in the domain.
EMC Navisphere Manager	The EMC Navisphere Manager application.
Event Monitor	EMC Navisphere Event Monitor.
FC-Series storage system	An FC4700, FC4700-2, FC4500, FC5600/5700, FC5200/5300 or FC5000 series Fibre Channel storage system.
FLARE™ software	Formerly known as Core or Base software.
HBA	Host bus adapter
Host Agent	The EMC Navisphere Agent that runs on a host system.
Initiator	A port on a NIC or iSCSI HBA that issues I/O requests to a target in the storage system. NICs and HBAs are installed on the environment's servers. The initiator is identified by a unique initiator name.
iSCSI	Internet SCSI (small computer systems interface)

Term	Meaning
iSCSI initiator	An iSCSI device that initiates requests to iSCSI targets. For example, a server that contains wither NICs or iSCSI HBAs could be an iSCSI initiator.
iSCSI target	An iSCSI device that responds to requests from an iSCSI initiator.
managed Agent	A Host Agent or SP Agent that you selected to manage.
managed Host	A Host Agent that you selected to manage.
managed storage system	A storage system managed by Navisphere Manager.
Manager	See EMC Navisphere Manager
metaLUN	A LUN that has been expanded from its original size by the addition of another LUN or by an entire physical disk. All metaLUN components (the original LUN and any added LUNs) are viewable from Navisphere, but you cannot divide a metaLUN into its original LUNs.
multi-domain environment	A group of domains that you can manage from a single instance of Navisphere Manager.
NAS	Network attached storage.
node	Usually a server (HBA connection) or storage system (SP connection). In the context of the web-based Manager, a node is an SP in a current storage system. The master node is the node that distributes the directory information to all storage systems in the domain. The master node always has an up-to-date copy of the directory.
Non-RAID Group storage system	A storage system whose SPs are running FLARE (formerly known as Core or Base) software without RAID Group functionality.
pre-FC4700 storage system	An FC4500, FC5300, FC5500, FC5700 or C-Series storage system.
RAID Group storage system	A storage system whose SPs are running FLARE software with RAID group functionality.
Server	A host connected (directly or through a switch) to the front-end or data ports on the storage-system SPs. The front-end connection can be Fibre Channel or iSCSI. A server can also be a management station.
SAN	Storage area network.
shared storage system	A storage system with the EMC Access Logix™ option, which provides data access control (Storage Groups) and configuration access control. A shared storage system is always a RAID Group storage system.
single-SP storage systems	see AX-Series storage system

Term	Meaning
software packages	Manager bundles storage-system software into <i>packages</i> . Each package has a name, revision, and attributes (for example, whether a package is revertible or needs committing).
SP	Storage Processor
SP Agent	The Navisphere Agent that runs in an SP (CX-Series and FC-Series storage systems).
storage application server	A storage system that has Navisphere Manager software installed in addition to Storage Management Server software. Using Manager, you can manage any storage system in the domain (depending on the type of user account).
storage management server	A storage system or SP that runs the Storage Management Server software.
Storage Group	Storage Group is a collection of one or more LUNs that you select, and to which you can connect one or more servers.
system	See storage management server .
Target	A storage system SP port that accepts and responds to requests from an initiator. In CLARiiON iSCSI systems, the targets are front-end or data ports on SPs.
unshared storage system	A storage system without the EMC Access Logix™ option.
Windows	Windows NT® or Windows® 2000, Windows Server™ 2003.

About EMC Navisphere software

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

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

Navisphere Storage-System Initialization Utility

The Navisphere Storage System Initialization Utility is supported for both AX-Series and CX-Series storage systems. You must initialize the storage system in order to register the server HBAs with the storage system and have access to the storage system's data.

For CX-Series storage systems, use the utility to discover storage systems, and set network parameters (IP address, subnet mask and default gateway). An authorized service provider must install and run the Initialization Utility.

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

Important The Navisphere Initialization Utility is supported only on the CX300, CX500, CX700, and AX-Series storage systems. The CX-Series systems must be running FLARE OE 02.16.xxx.5.yyy or higher. The AX100-Series systems must be running FLARE OE 02.16.xxx.5.yyy or higher. The AX150-Series systems must be running FLARE OE 02.20.xxx.5.yyy or higher.

Navisphere Host and SP Agents

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

Host Agents reside on attached hosts. SP Agents reside on the storage processors (SPs) of the most current storage systems (such as the CX-Series, FC-Series, and AX-Series), and are installed at the factory.

Navisphere Server Utility

Similar to the Host Agent, the Navisphere Server Utility registers the server's HBA (host bus adapter) with the attached AX-Series or CX-Series storage system; however, unlike the Host Agent, the Server Utility does not

- ◆ send LUN mapping information back to a CX-Series storage system. This information will not display in the Manager UI or CLI. To see LUN mapping information in the Manager UI or CLI for a CX-Series storage system, you must use the Host Agent instead of the Server Utility.

LUN mapping information is sent to AX-Series storage systems. In Navisphere Manager, this information displays in the Manager UI or with the CLI **-lunmapinfo** command. In Navisphere Express, this information displays in the **Manage Virtual Disks** page.

- ◆ automatically retrieve information from storage systems, and forward it to Navisphere Manager or Navisphere Express. With the Server Utility, you must manually update the information by starting the utility or you can create a script to run the utility.

For Windows servers running version 6.20 or higher, you can automatically update server information if the **Registration Service** feature remains enabled after installing the Server Utility.

- ◆ require network connectivity to the storage system.

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

However, with the Server Utility you can perform the following functions:

- ◆ Windows servers with Microsoft iSCSI initiators can use the utility to configure iSCSI connections - log on, log off, and remove an iSCSI target - and to configure mutual CHAP.
- ◆ AX-Series storage systems running Navisphere Express can use the utility to manage snapshots. If you are managing your AX-Series storage system with Navisphere Manager or if you have a CX-Series storage system, you must use SnapView and the admsnap utility to manage snapshots.

Important The Server Utility is supported on the AX-Series and CX-Series storage systems. The CX-Series systems must be running FLARE OE 02.16.xxx.5.yyy or higher. The AX-Series systems can run any supported revision of FLARE; however, not all features are supported in all revisions of the Server Utility. Refer to the Server Utility's online help for information on which features are available.

You cannot install the Host Agent and the Navisphere Server Utility on the same server. However, you can install them on different servers that are connected to the same storage system.

Navisphere Manager

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

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

Manager 6.X is a web-based user interface that lets you securely manage CLARiiON storage systems locally on the same LAN or remotely over the Internet, using a common browser. Manager 6.X resides on a CX-Series, FC-Series or AX-Series storage system, or a Windows Server™ 2003, Windows® 2000, or Windows NT® server that is running the Storage Management Server software, and is downloaded to the browser when the Storage Management Server software is accessed.

For an AX-Series storage system, you must complete an upgrade from Navisphere Express to Navisphere Manager. Reference the *AX100-Series Upgrading a Storage System to Navisphere Manager* document and the *EMC Navisphere Manager Administrator's Guide*.

Navisphere Integrator

Integrator allows Navisphere managed storage systems to be discovered and monitored through the leading enterprise management platforms (EMP). Integrator also enables the launch of Manager from the EMP. Integrator supports the following EMPs:

- ◆ HP OpenView® management software
- ◆ Tivoli NetView® management software
- ◆ CA Unicenter™ TNG Framework

Navisphere Integrator is installed and executed on a Windows 2000 or Windows NT server that also hosts the EMP.

Navisphere Storage Management Server

The Storage Management Server software is provided with Manager 6.X, and is installed and executed on each SP in a CX-Series or FC-Series storage system, or a Windows Server 2003, Windows 2000, or Windows NT server. A storage system with this software installed is a storage management server. A server with this software installed is called a server portal. All CX-Series (excluding the CX200 and CX300, AX-Series, and currently shipping FC4700-Series storage systems leave the factory with the Storage Management Server software installed on the SPs.

The CX200 and CX300 ship with the Storage Management Server Base software installed on the SPs.

The Storage Management Server software performs the following functions:

- ◆ Receives and responds to requests from Navisphere Manager 6.X
- ◆ Forwards requests to the local SP Agents for processing
- ◆ Forwards status and configuration updates to Navisphere Manager 6.X
- ◆ Replicates user and domain information to all storage systems in domain
- ◆ Authenticates user logins and authorizes user requests
- ◆ Logs all user logins and actions

A system with Storage Management Server software installed can also act as a portal for managing storage systems that do not have the Storage Management Server software installed (such as FC5300 and FC4500 storage systems). As a portal, the system forwards requests to the SP or Host Agent that is managing the storage system, and also collects information from the SP or Host Agent and relays that information to Manager 6.X for display updates.

Navisphere CLI

The CLI complements or can be used as an alternative to Manager. It provides a command line interface for storage-system management, including storage provisioning, status and configuration information retrieval, and control. You can use the CLI to automate management functions through shell scripts and batch files. CLI commands for many functions are server-based and are provided with the Host

Agent. The remaining CLI commands are web-based and are provided with software that runs in the storage-system SPs. Navisphere CLI includes the Secure (on supported operating systems), Classic, and Java CLI functionalities

The Navisphere CLI is supported on CX-Series, FC-Series, and AX-Series storage systems.

For an AX-Series system, you must install the Navisphere Manager Enabler in order for Navisphere CLI to be operational. Reference the *AX100-Series Upgrading a Storage System to Navisphere Manager* document. Once you install the Navisphere Manager Enabler, Navisphere Express is no longer operational. You cannot uninstall the Navisphere Manager Enabler once you install it. You cannot use Navisphere Express and Navisphere CLI to co-manage an AX-series storage system.

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

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

The trade-offs between Manager and the CLI are as follows

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

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

Navisphere Analyzer

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

Analyzer 6.X requires that Navisphere Manager 6.X be installed. Like Manager, Analyzer resides on a CX-Series or FC-Series storage system or a Windows NT® or Windows® 2000 server that is running the Storage Management Server software, and is downloaded to the browser when the Storage Management Server software is accessed.

Navisphere environments

You start Manager 6.X in a supported browser on a Windows or UNIX® host by entering the IP address of an SP in the CX-Series or FC-Series storage system with Manager installed on it. The Manager user interface (UI) is downloaded to the Windows or UNIX host and it runs within the browser. All requests from the UI go to the Storage Management Server software on the storage systems, which communicates with the SP Agent. The SP Agent, in turn, communicates with the FLARE or Access Logix™ software on the SPs.

With Navisphere 6.X, you can set up a portal system to manage systems that do not have Storage Management Server software installed (such as FC5300, FC4500, and NAS systems). For more information on configuring and using portals, refer to Chapter 17, "Managing Legacy Storage Systems and NAS Systems".

Storage-system configuration and management

Before you can configure or manage shared storage systems with Manager 6.X, you need to set up the Navisphere environment.

To set up Navisphere security and domains

1. Install and start Navisphere Manager.

To correctly manage the systems in the domain, Navisphere Manager must be at the same revision or higher than the revision of the FLARE Operating Environment running on the storage system.

2. Initialize security (refer to *EMC Navisphere Security Administrator's Guide*).
3. If you are setting up an iSCSI storage system, set up CHAP authentication as described in Chapter 16, "Managing iSCSI (Internet SCSI) Storage Systems".
4. Set up domains/multi domains, and user accounts (refer to the *EMC Navisphere Security Administrator's Guide*).

To install optional software on a CX-series or FC-series storage system

Install SnapView™, MirrorView™ or SAN Copy™ software on the CX-Series or FC-Series storage system. (Only a service provider should perform these install procedures).

For a pre-FC4700 storage system - Configure the Host Agent on the server (Chapter 3).

To configure a shared storage system with Manager

1. Set storage-system general properties, memory, and cache properties (Chapter 7).
2. If you want fair access to the storage system, set its host properties (Chapter 7).
3. Create RAID Groups and LUNs in the RAID Groups (Chapter 8).
4. Enable data access control for the storage system (Chapter 9).

5. Create Storage Groups and connect each server to its Storage Group (Chapter 9).
6. Make the LUNs available to the server's operating system. (refer to the operating system documentation).
7. For storage systems with the SnapView option, incremental SAN Copy option, or MirrorView / A option, set up the reserved LUN pool (see Chapter 11).
8. After you have configured all the storage systems connected to the configuration server, you can physically connect other servers to the storage system, or power up the other servers connected to the storage system.

To manage a shared storage system with Manager

1. Monitor storage-system operation and the failover software operation so you can recover from any failures (Chapter 12).
2. For a CX-Series or FC4700-Series storage system with the SnapView option, if you want to learn to use the SnapView option, refer to the *EMC SnapView for EMC Navisphere Administrator's Guide*.
3. For a CX-Series or FC4700-Series with the MirrorView option, if you want to learn to use the MirrorView option, refer to the *EMC MirrorView for EMC Navisphere Administrator's Guide*.
4. For a CX-Series or FC4700-Series with the SAN Copy option, if you want to learn to use the SAN Copy option, refer to the *EMC SAN Copy for EMC Navisphere Administrator's Guide*.
5. To manage an iSCSI storage system, refer to Chapter 16, "Managing iSCSI (Internet SCSI) Storage Systems".
6. Reconfigure Storage Groups, LUNs, or RAID Groups, if desired (Chapter 13).
7. Add a server to or remove one from a shared storage system, if desired (Chapter 13).
8. Reconfigure storage-system hardware, if desired (Chapter 13).

Preparing to Use and Starting Manager

This chapter describes

- ◆ Preparing to use Manager.....2-2
- ◆ Starting a Manager session.....2-3

Preparing to use Manager

Before you use Navisphere Manager 6.X to configure your storage, you must complete the following tasks:

Task	Described in
Install Navisphere Manager	<i>EMC Navisphere Manager Installation Sheet</i>
Confirm the revision of the JRE, and install a new revision if needed	<i>EMC Navisphere Manager and Manager Base Release Notes, and EMC Navisphere Manager Administrator's Guide, page 2-3.</i>
Initialize security on the storage system.	<i>EMC Navisphere Security Administrator's Guide</i>

The following tasks are optional depending on the tools and the language you will use to manage your storage systems:

Task	Described in
If you will be using Navisphere CLI to configure storage systems and you want to restrict access to these storage systems, do the following: - For CX-Series and FC4700-Series storage systems, add a privileged user to the SP Agent configuration file. - For pre-FC4700 storage systems, add a privileged user to the Host Agent configuration file. If the SP Agent or Host Agent configuration files do not include any privileged users, anyone who can log in to the management station can configure the storage system.	This document. Refer to "Adding privileged users to the SP agent configuration file" on page 7-20, and "Adding privileged users" on page 6-3.
Select a new language for the user interface and online help.	This document. Refer to the section, "Selecting a new language for Navisphere Manager" on page 2-5.

Installing a new version of the JRE

Please refer to the *Environment and System Requirements* section of the Release Notes for the supported versions of the JRE for this release.

To install the correct JRE, you may need to access the archive section of the Sun website, as follows:

1. At the Sun website, proceed to the Java directory. Click the **Products and APIs** link.
2. Under **Related Links**, click the **Download Archives/EOL** link.
3. Select the correct JRE version to download.

You can also use the following link to download the correct version of the JRE:

<http://java.sun.com/downloads/index.html>

Starting a Manager session

Use Navisphere Manager to create or modify domains, multi-domains and user accounts.

Before starting a management session, make sure that all storage systems you want to manage with this session are powered up.

1. Log in to a computer that is running a supported browser and Java Runtime Environment (JRE).

If you do not have a supported version of the JRE installed, on startup you may be directed to a version greater than that supported by Navisphere 6.X. See "Installing a new version of the JRE" on page 2-3.

2. Start the browser.
3. In the browser, enter the IP address of a storage system or Windows portal that has Manager installed.

The software displays a dialog box that requests your username, password, and scope (local or global) under which you want to log in. If you know the username and password of a valid account, go to step 4.

If no global administrator account exists, the software displays the message `Security not initialized` and asks if you want to initialize it, that is, if you want to create a global administrator account. Continue with this step.

- a. Answer **Yes**.

The software asks for a username.

- b. Define a username (1 to 32 letters (case sensitive), numbers, and/or underscores, must start with a letter).

The software asks for a password.

- c. For **Password**, enter 1 to 32 characters, including letters (case sensitive) and/or numbers.
- d. sensitive) and/or numbers.

Important – If the person who manages a storage-system installation cannot log in as a global administrator (perhaps because he or she forgot the password), then global management of the domain will be impossible. Any local administrators and managers retain their local management privileges. However, EMC support will need to recreate a global administrative account. So you should make sure that people who will manage the system keep a good record of the username and password.

- e. Click **OK**.
- f. From the **File** menu, select **Log In**.

4. Enter the username and password for the account and select the scope — **Global** or **Local**.

The Main window opens with one Enterprise Storage dialog box.

- In the **Storage** tree, Manager places a storage-system icon for every storage system in this domain.
- In the **Hosts** tree, Manager places a host icon for each server connected to any storage system in this domain.
- In the **Monitors** tree, Manager places a monitor icon for every monitored storage system in the domain, as well as for storage systems that do not belong to the domain, but are physically connected to storage system that do belong. Examples are legacy storage systems and SAN Copy destination storage systems.

Selecting a new language for Navisphere Manager

If you have language packs, other than English, installed on the same storage system with the Navisphere Manager UI (User Interface), you can select a new language for the user interface and online help for all installed applications.

If you do not upgrade the associated language pack when you install a new Navisphere 6.X software component, any new text in the software component will display in English, not in the selected language.

1. On the **Tools** menu, click **Select Language** to open the **Select Language** dialog box.
2. In **Select Language** choose a new language from the list of installed languages.
3. Click **OK** to save the change.
4. Re-start the web browser in order for the new language to take

What next?

- ◆ To learn about storage-system trees and the Main window, go to Appendix A, "Using Trees and the Main Window".
- ◆ To configure portals for managing legacy storage systems, go to Chapter 17, "Managing Legacy Storage Systems and NAS Systems".

This chapter describes how to install software on storage systems supported by the Software Installation wizard, and on pre-FC4700 storage systems. It also describes other software operations such as commit and revert, and how to install and upgrade disk firmware on all storage systems.

Topics include:

- ◆ Storage-system software installation introduction 3-2
- ◆ Software installation wizard overview 3-3
- ◆ Committing (finalizing) the software installation.....3-11
- ◆ Reverting back to the previous software version.....3-12
- ◆ Displaying software status for CX-series and FC4700-series storage systems.....3-12
- ◆ Using the newly installed software packages.....3-13
- ◆ Installing and upgrading software on pre-FC4700 storage systems 3-14

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

Storage-system software installation introduction

For CX-Series, FC4700-Series, and AX-Series storage systems, Navisphere Manager lets you perform the following four basic software installation operations:

- ◆ Install or upgrade storage-system software.
- ◆ Commit to finalize the software installation (refer to "Committing (finalizing) the software installation" on page 3-11).
- ◆ Revert back to the previous version of the software (refer to "Reverting back to the previous software version" on page 3-12).

For pre-FC4700 storage systems, Navisphere Manager lets you install or upgrade a storage system's FLARE software.

The procedure you use to install or upgrade storage-system software depends on whether you are installing the software on a pre-FC4700 storage system or on a storage system supported by the Storage System Software Installation wizard.

- ◆ For storage systems supported by the wizard, refer to "Software installation wizard overview" on page 3-3.

For storage systems not supported by the wizard, only a service provider should perform the installation procedure. Refer to the Navisphere Manager release notes for a list of supported storage systems.

- ◆ For pre-FC4700 storage systems, refer to "Installing and upgrading software on pre-FC4700 storage systems" on page 3-14.

Software installation wizard overview

The Software Installation Wizard (also known as the Smart Installation Wizard) helps you install the following software on EMC® CX-Series and AX-Series storage systems running Navisphere Manager:

- ◆ FLARE™ OE 02.07.xxx.5.yyy or higher bundle
- ◆ Navisphere® Manager Enabler
- ◆ Access Logix™ Enabler
- ◆ Analyzer Enabler
- ◆ MirrorView/S Enabler (synchronous MirrorView)
- ◆ MirrorView/A Enabler (asynchronous MirrorView)
- ◆ SAN Copy™ Enabler/SnapView™ Enabler

It also describes how to install Navisphere Management Server and Navisphere UIs for Microsoft Windows® hosts.

EMC does not currently support the wizard on all storage systems. Only a service provider should perform the installation procedure on unsupported storage systems. Refer to the Navisphere Manager release notes for a list of supported storage systems.

All storage systems ship from the factory with a version of FLARE OE installed. Some storage systems ship from the factory with some enablers installed, as listed in Table 3-1. To use the enabler interfaces, they must be enabled on the storage system to which you point your browser.

Table 3-1 Enablers for AX-series and CX-series storage systems

Storage system	Enablers installed on storage system prior to shipment from factory						
	Access Logix	Manager	MirrorView/S	MirrorView/A	SAN Copy	SnapView	Analyzer
AX-Series	Yes	Yes	N/A	N/A	No	Yes	N/A
CX200	No	No	N/A	N/A	N/A	N/A	No
CX300	Yes	Yes	N/A	N/A	N/A	No	No
CX400	No	No	No	No	No	No	No
CX500	Yes	Yes	No	No	No	No	No
CX600	No	No	No	No	No	No	No
CX700	Yes	Yes	No	No	No	No	No

Starting the software installation wizard

You should not initiate any configuration changes such as binding new LUNs or expanding metaLUNs while a software installation is in progress. Such operations will be rejected and will be delayed until Navisphere completes the software upgrade and until you commit the newly installed software.

1. Right-click the storage system on which you want to install the software and click **Software Operations > Software Installation Wizard**.

The wizard lets you install software on only one storage system at a time. Do *not* select multiple storage systems.

2. In the **Welcome** screen, read the wizard summary, and click **Next** to open the **Non-Disruptive Upgrade Delay** screen (storage system with two active SPs) or the **Software Package Selection** screen.

Entering non-disruptive upgrade delay

This screen will *not* be displayed if you have a single SP storage system or a dual SP storage system with only one active SP.

Important If the information in this screen is incorrect, you may lose access to your data for the duration of the software installation process.

Use this wizard screen to enter a non-disruptive upgrade delay. This allows attached servers enough time after SP B has finished rebooting to rescan and mark paths available to SP B before SP A reboots. We recommend that you keep the delay set to the 360 seconds (default value).

For servers running VERITAS Volume Manager DMP, enter the maximum path scan interval of the attached servers. The minimum scan interval is 360 seconds; the maximum is 1140 seconds. If you enter a value less than the minimum, the value will default to the 360 second minimum. You cannot enter a value greater than the maximum value. Manager automatically adds 60 seconds to this poll rate to insure that the primary SP does not reboot until the secondary SP is fully operational.

Manager saves the scan interval for the duration of this session of the wizard. For more details on scan interval refer to the VERITAS documentation.

Click **Next** to open the **Software Package Selection** screen.

Selecting software

Use this wizard screen to select the software you want to install on the storage system.

Currently Installed Software lists and describes any software that is already installed and running on the storage system.

1. Click **Browse** to open the **Open File** dialog box. This dialog box lets you select one or more software packages to install, lets you filter the files so that you see only those files of the type you want to install, and lets you select files from different directories or different CDs.
2. In the **Files of Type** drop-down menu, select one of the following:
 - **.lst, .pbu, .ena** where
 - .lst** file is the software bundle list file
 - .pbu** is a zipped software bundle file
 - .ena** is a software enabler file
 - **.lst, .pbu, .ena, .upf** where
 - .upf** is an upgrade package file
 - **All Files**

You do not need to manually unzip the **.pbu** file. If the storage system you are upgrading is running version 02.17.xxx or lower software, any **.pbu** files will be unpacked (extracted) to the client. For storage systems running version 02.19.xxx and higher software, the **.pbu** files will be extracted to the storage system.

Browse automatically inserts the double quotation marks and spaces needed for multiple file selection. You must select at least one software package.

The wizard places any selected software packages in the **Software to Install - Selected Files** list.

3. To remove a package from the **Selected Files** list, select the package, and then click **Remove**, or to remove all packages from the list, click **Clear All**.

The wizard saves the list of software packages for the duration of this wizard session.

4. Click **Next** to open the **Pre-install Validation** screen.

Pre-install validation

If the version of the storage system software that you are installing is later than the wizard software, the wizard upgrades itself to the latest rules.

This screen describes and performs a series of checks to verify that your current storage-system settings meet the installation or upgrade criteria. First, the wizard performs a version compatibility check to ensure that the storage-system software is capable of performing the remaining checks. If the version check succeeds, the wizard performs the remaining checks in parallel. Validation checks are described in Table 3-2.

Table 3-2 Validation checks

Rule	Description
Version Compatibility	Verifies that the wizard is compatible with storage-system software. If this check fails, the wizard will not perform any other validation check.
Special Conditions Warning	Reminds you to verify the following: • All attached servers are running failover software • All servers have at least one working path to each storage processor (SP). • You have stopped I/O to all attached VMware® ESX Servers™ and AIX servers. This rule will always result in a warning.
Acceptable processor Utilization	Verifies the SP Utilization. The SP utilization must be less than or equal to the maximum of 50%. If this rule fails, you must wait until the storage system is not as busy.
Redundant SPs	Verifies the availability of the SPs. For dual-SP storage systems, where one SP is not operating normally, Navisphere displays an error icon and prevents the installation until the problem is fixed. If the storage system has only one SP, Navisphere displays a warning icon.
Host Connectivity	Verifies that there is a connection path from each attached host to each storage system SP. CX200LC (single SP) and AX-Series storage systems will always return a warning.
No Trespassed LUNs	Verifies that there are no trespassed LUNs on the storage system. If LUNs are trespassed, a Navisphere displays warning icon. Note LUNs are recognized as trespassed only if they are in a Storage Group.

Table 3-2 Validation checks

Rule	Description
No Transitions	Verifies that no RAID Groups are expanding, no LUNs are binding or expanding, and that no disks are equalizing or rebuilding on the storage system.
No System Faults	Verifies that no faults exist on the storage system. The only fault that is acceptable is a charging SPS.
ATA Drive Rule	Verifies that the storage system has no ATA drives. If ATA drives are present, a service provider must perform the software installation. Note This rule is run only on storage systems running FLARE operating environment version 02.07.xxx.5.yyy or earlier. If your storage system is running a later version, this rule will not be run.
All Packages Committed	Verifies that all packages installed on the storage system are committed. A fix procedure is available for this rule.
Active Sessions Rule	Verifies that there are no active SnapView, Incremental SAN Copy, or MirrorView/Asynchronous sessions. You must stop any sessions before performing the software installation. Note This rule runs only on storage systems running FLARE operating environment version 02.16.xxx.5.yyy or later. If your storage system is running an earlier version, this rule does not run.
Acceptable I/O Rate	Verifies that the disk I/O rate is acceptable for the install operation.
Disruptive Upgrade Warning (for AX Platforms)	Warns that a software upgrade on a AX-Series storage system will always be disruptive - both SPs will reboot at the same time.
Statistics Logging Disabled	Verifies that statistics logging is disabled on the storage system. If not, the Fix button becomes active.
No Private MetaLUNs	Determines if any private metaLUNs (configured as write intent log or clone private LUNs) will interfere with the software installation. Upgrades to version 02.18.xxx.5.yyy or higher software will fail.
No Layered Applications	Verifies that layered drivers are not enabled on the storage system. If the storage system you selected is running the FLARE operating environment, version 02.07.x.y.z or earlier, and the software detects layered drivers on the storage system, a service provider must perform the software installation. Note This rule runs only on storage systems running FLARE operating environment version 02.07.xxx.5.yyy or earlier. If your storage system is running a later version, this rule does run.

After the wizard performs the validation checks, a status icon indicating if the rule is met or if action is required displays next to each rule. If you select the icon, status information displays in the **Rule Check Information** section. If action is required, a fix icon may display, which would initiate a fix when you select the icon. If a fix

icon does not display, refer to the **Rule Check Information** for details on how to fix the problem.

Important If a failed icon displays, you cannot continue with the installation until you fix the problem. If you click **Next** without performing the required fix, the wizard will restart the validation checks. If you fix the failure and click **Next**, the wizard will continue to the next screen.

The wizard displays one of the following icons:

Table 3-3 Validation check icons

Icon	Description
	Validation check in progress.
	Validation check in progress.
	Rule validation check succeeded with a warning. A fix is not required.
	Rule validation check failed and a fix is required.
	When selected, initiates a fix for the rule that failed. Note A fix is available only for the All Packages Committed rule.

5. Click **Next** to open the **Summary** screen.

Viewing installation summary

The **Summary** screen displays the following summary information:

- ◆ A list of software packages that will be installed on the storage system.
- ◆ A summary of installation settings, for example non-disruptive upgrade delay.
- ◆ The overall install type (for example, **Non-disruptive**) for this installation.

Click **Next** to start the installation and open the **Installation Status** screen.


If the installation fails, the wizard displays a Warning message that describes the reason for the failure.

Viewing and monitoring installation status

The **Installation Status** screen displays a message indicating the successful start of the software installation.

To monitor installation progress, click **Finish**. The **Software Operation Progress History** dialog box opens.

You can also open this dialog box by right-clicking a storage-system icon, and then clicking **Software Operations > Software Operation Progress History**.



This dialog box displays the software operation and software packages selected, and while the operation is ongoing, tracks its progress. The dialog box lists the major steps required for the operation, adds a checkmark to each completed step, and when the operation completes, displays the result of the operation.

If possible, when the installation or upgrade is complete, the software restores the cache to its original size, and asks you to verify that this was done correctly.

What next?

If you are ready to finalize the installation of FLARE and Access Logix software, go to the next section, "Committing (finalizing) the software installation".

Committing (finalizing) the software installation

You cannot create any new RAID Groups, or bind any new LUNs until you commit FLARE and Access Logix software.

New features will not be available until the commit is completed. When you are ready to use the software in a production environment, you must commit the software to finalize the installation.

The write cache is flushed as part of the commit operation. This can be a time consuming operation, and Navisphere Manager may time out waiting for a response from the storage system. If this happens, Navisphere Manager will display one or both SPs as unmanaged. However, even though the SPs are unmanaged, they are still servicing I/O in this state. Once the commit is complete, Manager displays the state of the SPs as normal.



CAUTION

Once you commit a software package, you cannot revert it.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose software you want to commit, right-click, and select **Properties**.
2. Use the **Software** tab to select and commit the software packages.

These packages will have an **Active (commit required)** status.

What next?

If you need to return to a previous revision software package, go to the next section, "Reverting back to the previous software version".

Reverting back to the previous software version

You may want to install new software packages and use them on a trial basis before you commit or finalize the installation. If the new software does not behave as expected, **Revert** lets you return to the previous revision software package, if one exists.

The software can be reverted only if:

- ◆ A previous version of the software is installed on the system
 - ◆ The software package has not been committed
1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose software you want to revert, right-click, and then select **Properties**.
 2. Use the **Software** tab to select and revert the software package.

Displaying software status for CX-series and FC4700-series storage systems

You can view the revision and current status of all software packages installed on CX-Series or FC4700-Series storage systems, as well as the dependencies of one software package on another.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose software status you want to display, right-click, and then select **Properties**.
2. Use the **Software** tab to view the status of all installed software and any software dependencies.

Using the newly installed software packages

In order to successfully use the newly installed software applications, such as Navisphere Manager, SnapView, SAN Copy, MirrorView or Analyzer, EMC recommends that you do the following:

1. Close the instance of the browser that is pointing at the newly updated storage system.
2. Clear the JRE and browser cache.
3. Open a new browser and enter the IP address of the storage system with the updated Navisphere software.

If you are upgrading FLARE software on multiple storage systems, you may want to perform two download operations: one for storage systems supporting warm reboot, and one for the remaining storage systems.

4. Click **Next** to start the download operation.

After the SPs reboot, you may have to restart the Host Agent on the servers connected to the storage system that received the new FLARE software.

Installing and upgrading software on pre-FC4700 storage systems

You can upgrade the FLARE software on a pre-FC4700 storage system. For a new revision of the FLARE software to take effect, the SPs in the storage system must be rebooted.

The FLARE software media for all pre-FC4700 storage systems may also include an upgrade to the SP programmable read-only memory (PROM) code. If PROM code is included, it is installed automatically with the FLARE software.

When you install FLARE software, the SP tries to copy it to reserved areas outside operating system control on several disks, which are called the database disks. Having multiple copies of code offers higher availability if a disk fails. The database disks for the different storage-system types are listed in Table 3-4.

Table 3-4 Database disks for different storage-system types

Storage-system type	Database disk IDs
C3x00 series	A0, B0, C0, A3, A4
C2x000 series	A0, B0, C0, A3
C1900 series	A0, B0, C0, A1
C1000 series	A0, A1, A3, B0

When you install FLARE software, at least two of the database disks must be online, and ideally, all of them should be online. A disk is online if it is fully powered up and not faulted; that is, if **Current State** is Normal on its **Disk Properties** dialog box. If you try to power up the storage system without two of these disks in place, the powerup fails.

The file for the new FLARE software revision must be on a server that can be reached across a network from the server connected to the storage systems whose FLARE software you want to upgrade.

To upgrade FLARE software

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the system or systems on which you want to install or upgrade the FLARE software, right-click, and then select **Software Operations > Software Installation**.

All storage systems must be the same type.

2. Under **Firmware Options**, there will be two choices:
 - If the new FLARE software file is accessible from the management station, select the **File Accessed Locally** check box before you click **Browse** to locate and then select the new FLARE software file.
 - If the new FLARE software file is accessible on a server connected to a storage system, clear the **File Accessed Locally** check box, and then in **Filename(s)**, type the complete pathname for the new FLARE software file.

Under **Reboot Options**, select one of the following check boxes:

The SPs in a storage system must be rebooted for a new revision of FLARE software to take effect.

- **No Reboot** will not automatically reboot the storage system after downloading new FLARE software.
- **Warm Reboot**, if available, suspends all outstanding I/O to the storage system, starts the SPs after downloading the new software, and typically takes less than 30 seconds.
- **Hard Reboot** terminates all outstanding I/O to the storage system, restarts the SPs after downloading the new software, and typically takes about 1 to 2 minutes. Before selecting **Hard Reboot**, suspend all I/O to the storage system.

If you are upgrading FLARE software on multiple storage systems, you may want to perform two download operations: one for storage systems supporting warm reboot, and one for the remaining storage systems.

3. Click **Next** to start the download operation.

After the SPs reboot, you may have to restart the Host Agent on the servers connected to the storage system that received the new FLARE software.

What next? If you will be using failover software, proceed to the next chapter. To configure the Navisphere Host Agents on the storage-system server connected to the pre-FC4700 storage systems, continue to Chapter 6, "Configuring the Remote Host Agent".

Managing and Monitoring AX-Series Storage Systems

This chapter describes

- ◆ Managing AX-series systems with Navisphere Manager 4-2
- ◆ Navisphere Manager behavior with AX-series systems 4-4
- ◆ Managing and monitoring configurations 4-13

Managing AX-series systems with Navisphere Manager

The term AX-Series storage systems refers to both AX100-Series and AX150-Series storage systems.

The primary management tool for AX-Series storage systems is Navisphere Express. With Navisphere Express, each AX-Series storage system is in its own domain, and as a result, you can manage only one AX-Series storage system at a time. With Navisphere Manager you can place multiple AX-Series storage systems in the same domain or multi-domain environment and manage and monitor more than one storage system at a time. In addition, you can place both CX-Series and AX-Series storage systems in the same domain or multi-domain environment and manage and monitor both storage system types with Navisphere Manager.

Once Navisphere Manager is installed on the AX100 storage system, Navisphere Express is no longer operational. You *cannot* co-manage the storage system using Navisphere Manager and Navisphere Express, and you cannot install Navisphere Express on a system with Navisphere Manager.

When you use Navisphere Manager to manage AX-Series storage systems, some Navisphere Manager commands are limited or unavailable.

Managing Prerequisites

Before you use Navisphere Manager as the management tool for an AX-Series storage system, complete the following prerequisites:

AX100-Series prerequisites

- ◆ Install the FLARE Operating Environment software, version 02.19.xx.5.yyy or higher.
- ◆ Install the Navisphere Manager Enabler software.

AX150-Series prerequisites

- ◆ Install the Manager Enabler software on any AX150 storage system that you want to manage.

If you are adding the system to a domain or multi-domain environment, the environment must include an off-array server running version 02.20.xx.5.yy of the Navisphere UI.

AX-Series storage systems can be managed by Navisphere Manager when they are part of any of the following storage environment configurations — single-domain environment, multi-domain environment, and multi-system domain. refer to "Managing and monitoring configurations" on page 4-13.

Navisphere Manager behavior with AX-series systems

Since Manager's user interface must include functionality for all storage system types it supports, some Navisphere Manager commands and features are limited or unavailable for AX-Series storage systems. When using Navisphere Manager with AX-Series storage systems, the following information will be helpful.

AX-series to CX-series terminology differences

The following table lists and defines the differences between AX-Series and CX-Series storage system terminology.

AX-series term	CX-series term	Navisphere Manager definition
virtual disk	LUN	A grouping of one or more disks or disk partitions into one span of disk storage space. A LUN/virtual disk looks like an individual disk to the server's operating system. It has a RAID type and properties that define it.
disk pools	RAID Groups	A set of disks on which you bind one or more LUNs. Each LUN you bind on a RAID Group is distributed equally across the disks in the Group. Each RAID Group supports only one RAID type.
RAID Type	RAID Type	The RAID type of a RAID Group determines the type of redundancy, and therefore the data integrity of the LUNs bound on it. AX-Series storage systems support only the Hot Spare, RAID 5, and RAID 1/0 RAID types.
N/A	Storage Group	A collection of one or more LUNs or metaLUNs to which you connect one or more servers. These LUNs are accessible only to the server that is connected to the Storage Group.
N/A	reserved LUN pool	The reserved LUN pool works with replication software, such as SnapView, SAN Copy, MirrorView/A, and MirrorView/S to store data or information required to complete a replication task.
N/A	templates	Specifies the storage system events that you want to monitor and method of event notification. A template contains events, responses, and message formatting and can be mapped to one or more storage systems.

Status reporting differences


This section describes any differences in the way Navisphere Express and Navisphere Manager report the status of storage devices.

Monitoring events

Event descriptions that appear in the **Events** dialog box in Navisphere Manager are not equivalent to those that appear in the **View Events** page in Navisphere Express.

AX-series tree structure

The Navisphere Manager tree structure for the AX-Series storage system displays some items that are specific to the AX-Series storage systems and not other storage system types. The following sample tree structure shows an AX-Series storage system with one SP and one NVRAM memory card. Your storage system might be configured with two SPs and no NVRAM card.



Right-click menu options for AX-series storage systems

When using Navisphere Manager with an AX-Series storage system, the following right-click menu options are available.

Monitoring only

When using Navisphere Manager to just monitor the AX-Series storage systems, Navisphere Express is still the management tool. In this configuration, you cannot expand the storage system icon in the Navisphere Manager tree structure.

Component	Tree	Menu option	Use to
Storage system	Storage, Hosts	Navisphere Express	Launch the Navisphere Express Management UI
		Faults	For systems running version 02.19.xx.5.yyy or higher of the AX-Series storage-system integrated management software, Manager displays the Navisphere Express Needs Attention messages in the Fault Status dialog box. For systems running earlier versions of the software, Navisphere displays standard error messages in the Fault Status dialog box.
		Shutdown	Shuts down the AX-Series storage system *.

* While the system is shutting down, you will lose access to all virtual disks in the storage system. The system will remain shut down until you power it on. This is a safe operation and you will not lose any data on the disks or in the storage system cache. Before you shut down the storage system, do the following:

1. Stop all applications running on servers connected to the storage system (stop data I/O from the servers to the storage system).
2. Save all unsaved data to disk. The steps required to complete this operation depend on the server's operating system. (Refer to your operating system documentation for the correct steps).

Managing and monitoring

Some options in the Navisphere Manager dialog boxes may be dimmed and unavailable to AX-Series storage systems. Refer to "Dialog box limitations" on page 4-11.

Component	Tree	Menu option	Use to
Storage system	Storage, Hosts	Create RAID Group	Create a RAID Group from selected disks (RAID Group storage systems only).
	Storage, Hosts	Bind LUN	Bind disks into logical units.
	Storage, Hosts	Create Storage Groups	Create a Storage Group on the storage system.
	Storage, Hosts	Faults	Display the fault status report for the storage system.
	Storage, Hosts	Disk Summary	Display a summary of the disk in the storage system.
	Storage, Hosts	Software Operations >	
		Software Installation Wizard	Upgrade storage system software
		Software Operation Progress History	Review the status or track the progress of software installation.
	Storage, Hosts	Connectivity Status	Display the status for each HBA connected to the storage system (Access Control must be enabled)
	Storage, Hosts	Update Now	Refresh the storage system status information displayed in the Enterprise Storage trees.
	Storage, Hosts	Properties	Display the properties of the storage system.
	Storage, Hosts	SnapView > (only if software is installed)	Refer to online help for SnapView.
Enclosure	Storage	Flash LEDs ON	Start flashing the LEDs (status lights) on the enclosure.
	Storage	Flash LEDs OFF	Stop flashing the LEDs (status lights) on the enclosure.
	Storage	Resume for MidPlane	Display properties for enclosure midplane.
Disk	Storage	Properties	Display the properties of the disk.
Fan	Storage	Properties	Display the state of the fan.

Component	Tree	Menu option	Use to
Power Supplies	Storage	Resume	Display the properties of the power supply.
	Storage	Properties	Display the current state of the power supplies.
SP	Storage	View Events	Display the event log for the storage processor (SP).
	Storage	Reset Statistics Logging	Not supported on AX-Series.
	Storage	File Transfer Manager	Transfer SP diagnostics log files to a local client.
	Storage	Resume	Display information stored in the resume prom for the selected SP.
	Storage	SP Collect	Generate diagnostic log files for the selected SP.
	Storage	Properties	Display the properties of the SP.
Standby Power Supplies	Storage	Properties	Display the properties of the SPS (standby power supplies).
	Storage	Set Test Time	Specify a day of the week and time of day for the application to conduct a battery test.
RAID Group	Storage	Destroy	Destroy the RAID Group and unbind its LUNs.
	Storage	Properties	Display the properties of the RAID Group.
	Storage	Bind LUN	Bind disks into logical units (LUNs) on a RAID Group.
LUN	Storage, Hosts	Unbind LUN	Destroy the LUN and unbind its disks.
	Storage, Hosts	Bring LUN Online	Bring LUNs that are in the Offline State back on line. These LUNs are identified in the tree structure as being "Offline".
	Storage, Hosts	Trespass	Transfer ownership of the LUN from the current SP to the peer SP. Not available for single SP storage systems.
	Storage, Hosts	Update Host Information	Not supported on AX-Series storage systems.

Component	Tree	Menu option	Use to
LUN (cont)	Storage, Hosts	Add to Storage Group	Add the LUN to a selected Storage Group.
	Storage, Hosts	Expand	Dynamically expand the capacity of a LUN by combining multiple LUNs into a single unit called a metaLUN. Only concatenation expansion is supported for AX-Series LUNs.
	Storage, Hosts	Properties	Display the properties of the LUN.
Storage Groups	Storage, Hosts	Destroy	Destroy a Storage Group on a storage system.
	Storage, Hosts	Select LUNs	Add LUNs to the Storage Group.
	Storage, Hosts	Connect Hosts	Connect hosts to a Storage Group.
	Storage, Hosts	Properties	Display the properties of a Storage Group.
Hosts	Storage, Hosts	Properties	Display the properties of the host.
	Storage, Hosts	Update Now	Refresh the host status information displayed in the Enterprise Storage trees.
	Storage, Hosts	Faults	Display the fault status report for the host.
	Storage, Hosts	Connectivity Status	Maintain connectivity to your storage after replacing a host HBA without having to disconnect the host from and then re-connect the host to Storage Groups.
Cache Memory Card (single SP systems only)	Storage	Initialize	Initialize the NVRAM card - displays a warning dialog box that describes the current condition of the memory card.
	Storage	Properties	Display properties of cache memory card.

Dialog box limitations

The following table lists the Navisphere Manager dialog boxes with options that are unavailable when managing AX-Series storage systems with Navisphere Manager.

Dialog box name	Unavailable options
Bind LUN	Element Size, Rebuild Priority, Verify Priority, Enable Read Cache, Enable Write Cache, Enable Auto Assign, Alignment Offset, No Initial Verify, LUN Size - TB, LUN Size - Block Count. RAID Type - only RAID 5 and RAID 1/0 are available
Configure Reserved LUN Pool	Add to SP B LUN Pool for single SP systems.
Create Initiator Record	Initiator type, Unit Serial Number, ArrayCommPath, Failover Mode
Create Raid Group	Only supports Hot Spare, RAID 5 and RAID 1/0 Raid Types.
Group Edit Initiators	Initiator Type, Unit serial Number, ArrayCommPath, Failover Mode.
Host Properties - Agent Tab	Not visible
Host Properties - Storage Tab	Update Host Side Data
LUN Properties - General Tab	Rebuild Priority, Verify Priority, Auto Assignment Enabled, Default Owner
LUN_Properties - Cache Tab	Read Cache Enables, Write Cache Enabled
LUN Properties - Prefetch Tab	All options
LUN Properties - Statistics Tab	All statistics are unavailable for AX-Series storage systems
LUN Properties - Hosts Tab	Update Host Information
Reserved LUN Pool Properties - SP B LUN Pool Tab	Not visible for single SP systems
SP Properties - Statistics Tab	All statistics are unavailable for AX-Series storage systems
SP Properties - ALPA Tab	Change button (cannot change ALPA settings)
SP Properties - Fibre Speed Tab	All options
Storage Group Properties - Storage Access Tab	Include connected option in Show Hosts drop down list
Storage System Properties - Cache Tab	All options
Storage System Properties - Memory Tab	All options
Storage System Properties - Storage Access Tab	Enforce Fair Access, Access Control Enabled
Expand Storage wizard - Specify Expansion Type	Stripe expansion

Dialog box name	Unavailable options
Expand Storage wizard - Specify MetaLUN Capacity	User Capacity: Blocks
Expand Storage wizard - Specify MetaLUN Settings	Default Owner, Expansion Rate, Element Size Multiplier, Alignment Offset, Enable Auto Assign
Software Installation Wizard	Only supports Disruptive NDU Does not execute Non-Disruptive Delay step
Template Properties - Paging Tab	All options
Events	none

Managing and monitoring configurations

This section describes the supported configurations for managing and monitoring AX-Series storage systems using Navisphere Manager.

Managing and monitoring in a single-domain environment

This procedure assumes that you have satisfied the prerequisites for managing the AX-Series storage system with Navisphere Manager, and that the AX100 storage systems will remain in their own domain and not be part of a multi-domain environment or a multi-system domain. An SP within the AX-Series storage system is the master of this domain.

1. Start Navisphere Manager by opening a web browser and entering the IP Address of an SP in the AX100 storage system.

The Navisphere Manager user interface opens.

2. Enter the username and password for the storage system.
3. Use Navisphere Manager to configure, manage and monitor storage.

Managing and monitoring in a multi-domain environment

If the AX-Series storage systems are located remotely, you may not want to include them in the local domain. You can add them as single-storage system domains to a multi-domain environment that includes a domain with a storage system or off-array server running the Manager UI. This procedure assumes that you have satisfied the prerequisites for managing the AX-Series storage system with Navisphere Manager, and you are adding each AX100 storage system as a *single domain* to a multi-domain environment.

When monitoring in a multi-domain environment, Navisphere Express is still the management tool for the AX-Series, and no AX-Series prerequisites are required.

With this configuration

- ◆ Unless manually modified, the storage systems retain their global and local user account credentials.

- ◆ Unless you create a user account on each storage system that matches an account on the local domain, you must use the login-on-demand feature to login to the storage systems

Managing and monitoring in a multi-system domain

An AX-Series storage system cannot be the master of a multi-system domain. Use a CX-Series storage system or an off-array server as the domain master.

This configuration provides consistent user account information for all storage systems in the local domain. Use it to simplify the administration of your storage enterprise. When you move an AX-Series storage system into your local domain, all global user accounts on the single AX storage system are replaced with the global administrator accounts from the local domain. Previous user accounts on the AX storage system no longer exist. This procedure assumes that you have satisfied the prerequisites and are adding AX-Series storage systems to the existing local domain. If you remove this system from the local domain, you must reinitialize it (assign new username and password).

When monitoring in a multi-system domain, Navisphere Express is still the management tool for the AX-Series, and no AX-Series prerequisites are required.

To add the storage system to a local domain, refer to the *EMC Navisphere Security — Domains, Multi-Domains, and User Accounts* manual, P/N 069001124.

Setting Storage System Failover Properties

This chapter describes

- ◆ Failover software values overview 5-2
- ◆ Setting failover values using the Failover Setup Wizard 5-3
- ◆ Setting failover values for initiators connected to a specific storage system 5-10
- ◆ Setting failover values when registering a new or existing initiator 5-10

Failover software values overview

To set failover option values using Navisphere Manager, your storage system must have Access Logix™ software installed. For storage systems without Access Logix, use Navisphere CLI.

To accommodate the failover software now available for CLARiiON storage systems, you must set specific options for each initiator. You must assign specific values to the **Initiator Type**, **Failover Mode** and **Array CommPath** options in order for the failover software to operate correctly.

For the correct failover settings, refer to the Installation Roadmap for your operating system, found on EMC Powerlink. Refer to the "Finding current information" section in the Preface.

- ◆ The **Initiator Type** is used for specific server OS configurations. Select **CLARiiON Open** for this field unless instructed otherwise.
- ◆ The **Failover Mode** option sets the behavior options based on the failover software installed on the server.
- ◆ The **Array CommPath** option enables or disables a communication path from the server to the storage system when LUN 0 is not configured on the storage system

Setting failover values using the Failover Setup Wizard

Using a series of dialog boxes, the **Failover Setup Wizard** helps you set failover values for all HBA initiators belonging to a specific server. The wizard updates the initiator records for this server in all storage systems connected to this server. Different server failover software (for example, ATF and Powerpath) requires different settings.

For the correct failover settings, refer to the Installation Roadmap for your operating system, found on EMC Powerlink. Refer to "Finding current information" in the Preface.

Before you use the **Failover Setup Wizard**, all HBA initiators belonging to this server must be registered (there must be an entry for each initiator in the storage system's server initiator table). To register any initiators, refer to "Registering HBAs and setting HBA initiator record values" on page 7-21.



CAUTION

If you enter incorrect values, the storage system could become unmanageable or unreachable by the server, or the failover software could stop operating correctly.

Starting and using the Failover Setup Wizard

1. To start the **Failover Setup Wizard**, click **Failover Setup Wizard**, on the **Tools** menu.

The **Start Wizard** dialog box opens, similar to the following:



2. In the **Start Wizard** dialog box, read the introduction to the wizard, and then click **Next** to open the **Select Host** dialog box.


The **Select Host** dialog box opens, similar to the following:



The **Select Host** dialog box displays a list of all currently managed hosts.

3. In **Host**, select the host that is running your failover software, and then click **Next** to open the **Select Storage Systems** dialog box.


The **Select Storage Systems** dialog box opens, similar to the following:



The **Select Storage Systems** dialog box displays a list of all storage systems that are currently connected to the selected host and have an entry for this host in its host initiator table.

4. Select those storage systems that you want to include in this failover configuration, and then click **Next**.

The **Specify Settings** dialog box opens, similar to the following:



5. In the **Specify Settings** dialog box, enter the correct values for **Initiator Type**, **Failover Mode**, and **Array CommPath**.

For the correct failover settings, refer to the Installation Roadmap for your operating system, found on EMC Powerlink. Refer to "Finding current information" in the Preface.




CAUTION

Entering incorrect values could cause the storage system to be unmanageable or unreachable by the host, and could cause the failover software to operate incorrectly.

6. Click **Next** to open the **Review and Commit Settings** dialog box.


The **Review and Commit Settings** dialog box opens, similar to the following:



7. In the **Review and Commit Settings** dialog box, review the failover configuration and all the settings.
8. If the settings are correct, click **Next** to continue; if the settings are incorrect, click **Back** until you return to the dialog box for which you need to re-enter values.

If you click **Next**, the wizard displays a confirmation dialog box. Click **Yes** to continue the operation, or click **No** to cancel it.

If you click **Yes**, the wizard opens the **Details from Operation** dialog box.



The **Details from Operation** dialog box displays a summary of the success of the operation for each storage system that you selected. If the operation failed on a storage system, the wizard displays the storage system name along with a description of the error.

9. For any failed storage systems, rerun the **Failover Setup Wizard**.
10. When the operation is successful for all storage systems, click **Finish** to close the wizard.

Setting failover values for initiators connected to a specific storage system

Navisphere Manager lets you edit or add storage system failover values for any or all of the HBA initiators that are connected to a storage system and displayed in the **Connectivity Status** dialog box for that storage system.

For the correct failover settings, refer to the Installation Roadmap for your operating system, found on EMC Powerlink. Refer to "Finding current information" in the Preface.



CAUTION

Entering incorrect values could cause the storage system to be unmanageable or unreachable by the host, and could cause the failover software to operate incorrectly.

1. In the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose failover properties you want to add or edit.
2. Right-click the storage-system icon, and click **Connectivity Status**.
3. In the **Connectivity Status** dialog box, click **Group Edit** to open the **Group Edit Initiators** dialog box.
4. Select the initiators whose **New Initiator Information** values you want to add or change, and then add or edit the values in **Initiator Type**, **ArrayCommPath** and **Failover Mode**.
5. Click **OK** to save the settings and close the dialog box.

Navisphere updates the initiator records for the selected initiators, and registers any unregistered initiators.

Setting failover values when registering a new or existing initiator

When you register a new or existing HBA initiator, Navisphere Manager lets you edit or add failover values to the initiator record. Refer to "Registering HBAs and setting HBA initiator record values" on page 7-21.

Configuring the Remote Host Agent

This chapter describes how to configure remote Host Agents for pre-FC4700 storage systems.

- ◆ Configuring Host Agents remotely (pre-FC4700 series) 6-2

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

Configuring Host Agents remotely (pre-FC4700 series)

If you have pre-FC4700 storage systems connected to a server, you must configure the Host Agent in the server so it can communicate with the FLARE software running in the storage-system SPs. Before you can configure a Host Agent using Navisphere Manager 6.X, the Host Agent configuration file must contain the IP address for *each* SP in a storage-system portal, or the IP address of the server portal.

What next?

- ◆ If no portal IP address exists, go to the section "Adding IP addresses to the agent configuration files" on page 17-8.
- ◆ If this user entry exists and the Host Agent configuration file does *not* include the line entry, **device auto auto** go to the section "Scanning for devices" on page 6-3.
- ◆ If this user entry exists and the Host Agent configuration file includes the line entry, **device auto auto**, go to the section "Updating parameters" on page 6-6.
- ◆ If you will also be using Navisphere CLI to configure storage systems, you must add a privileged user to the
 - Host Agent configuration file for pre-FC4700 storage systems (refer to the next section, "Adding privileged users").
 - SP Agent configuration file for CX-Series and FC4700-Series storage systems (refer to "Adding privileged users to the SP agent configuration file" on page 7-20).

Adding privileged users

Privileged users can configure the storage system, including binding and unbinding LUNs. When you add a privileged user, the system adds the user to the **agent.config** file. The easiest way to edit this file is with the remote agent configuration feature of Navisphere Manager. Before you can use Manager to add privileged users to the Host Agent configuration file, you must add the IP addresses for the portal to the **agent.config** file.

You must also add the IP addresses for the portal to the **agent.config** file in order to configure legacy storage systems, and to perform event monitoring actions on an attached server.

1. In the **Enterprise Storage** dialog box, click the **Hosts** Tab.
2. Right-click the icon for the server for which you want to display the agent properties, and click **Properties**.
3. Click the **Agent** tab.
4. Under **Privileged Users**, click on a blank line, and enter the name of the privileged user.
5. Click **OK** to add the privileged user and return to the **Agent** tab.
6. Click **Apply** in the **Agent** tab to save your changes and continue editing the agent configuration file, or click **OK** to save your changes and close the **Host Properties** dialog box.

You do not need to stop and start the Host Agent or click the Reload button for the changes to take effect.

Scanning for devices

Before the Host Agent can communicate with a storage system, you must add a communication channel (device entry) to the Host Agent configuration file.

If, when you were installing the server software, you edited the Host Agent configuration file to include the entry, **device auto auto**, you do not need to add communication channel device entries to the Host Agent configuration file. These are created dynamically each time you start the Host Agent.

1. From the **Enterprise Storage** dialog box, select the **Hosts** tab and navigate to the icon for the host whose Agent configuration file you want to configure, right-click, and then select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. In the **Host Properties** dialog box, click the **Agent** tab and scan for device entries by doing one of the following:

- Click **Auto Detect**; the Host Agent adds valid connection paths to the storage systems connected to the server.
- or
- Clear **Auto Detect**, and then click **Advanced**; the **Advanced Device Configuration** dialog box opens.

If the Host Agent configuration file was never edited to add device entries, the **Communications Channels** list is empty.

By clicking **Scan**, the Host Agent clears the **Communications Channels** list, scans the SCSI bus, and adds all connected storage systems it finds to the **Communication Channels** list.

or

- Click **Scan Bus**; the Host Agent adds valid connection paths to the storage systems connected to the server, and opens the **Scan SCSI Buses** dialog box.

Scan SCSI Buses lets you view specific information for all EMC CLARiiON storage devices and non-CLARiiON storage devices.

What Next?

- ◆ To make changes to the **Communications Channels** list, go to "Updating the communications channels list" on page 6-5.
- ◆ To update Host Agent parameters, go to "Updating parameters" on page 6-6.

Updating the communications channels list

You can add, delete or clear device entries from the **Communications Channels** list.

Configuring a pre-FC4700 storage system to allow one or more Agents to communicate with the same storage processor over multiple communication channels may lead to performance degradation and unpredictable results. We recommend not using such a configuration.

Adding devices

The remote Host Agent lets you add new devices to the **Communication Channels** list.

From the **Advanced Device Configuration** dialog box, click **Add Device** and do the following:

1. Enter the **OS device name**.
2. In **Storage System**, type the name of the storage system that you want the new device to manage.
3. In **Connection Type**, select the desired connection type.
4. Optionally, in **Comment**, type any comments pertaining to this device.

Deleting devices

When you delete a device, you remove it from the **Communication Channels** list. When you remove the device from the **Communication Channels** list, the device can no longer manage the storage system.

In the **Advanced Device Configuration** dialog box, select the device that you want to delete and click **Delete Device**.

Clearing devices

Clearing devices removes all the current devices from the **Communication Channels** list.

In the **Advanced Device Configuration** dialog box, click **Clear**. All connected devices are removed from the **Communication Channels** list.

Updating parameters

Updating parameters includes setting the serial line baud rate and the log entries to transfer. To update parameters, you must have privileges.

Serial Line Baud Rate - Lets you select the serial communication baud rates. Valid values are 9600, 19200, and 38400.

Log Entries to Transfer - Lets you select the log size to be transferred. Valid values are 100, 2048, 5000, and All.

Use the **Agent** tab in the **Host Properties** dialog box to set the parameters.

What next?

Continue to Chapter 7 to set storage system properties.

Setting Storage-System Properties

When you set up a storage system, you can view and change many of its general, memory, cache, storage access, and software properties or use the default values for these properties.

For security reasons, we recommend that you place all storage systems in a domain before you configure the them.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes

- ◆ Setting storage-system management control properties 7-2
- ◆ Setting storage-system general configuration properties 7-7
- ◆ Setting storage-system memory properties.....7-9
- ◆ Setting storage-system cache properties.....7-11
- ◆ Setting the fair access property (pre-FC4700 only) 7-16
- ◆ Setting the SP network and ALPA properties (CX-series or FC4700-series only).....7-17
- ◆ Setting the SP fibre speed property 7-19
- ◆ Adding privileged users to the SP agent configuration file 7-20
- ◆ Setting the battery test time.....7-20
- ◆ Registering HBAs and setting HBA initiator record values 7-21

Setting storage-system management control properties

Storage-system management control properties are available only for pre-FC4700 storage systems with Access Logix installed.

In order to use Navisphere Manager 6.X to manage pre-FC4700 storage systems the following must be true:

- ◆ The storage system must be part of a portal configuration. Refer to "Adding and configuring a portal" on page 17-5.
- ◆ You must add the IP address for both storage-system SPs, or the IP address of the server portal to the privileged user's list in the Host Agent configuration file. Refer to "Adding IP addresses to the agent configuration files" on page 17-8.
- ◆ You must have, at a minimum, a local user account, with the manager role, on a storage system or portal that has Manager 6.X installed.

This section describes:

- ◆ Management control
- ◆ Storage-system management control properties
- ◆ How to set storage-system management control properties

Management controls

Management control lets you restrict the servers that can send configuration commands to the storage system. We recommend that you use the storage system's management control option to give this privilege to one or two servers only.

A privileged user assigned the manager role can perform any configuration task, such as binding or unbinding LUNs from the management station. Management control lets you restrict the servers that can send configuration commands from this user to an attached storage system. Without management control, any server can send configuration commands to any connected storage system.

Management control is governed by a management login password that you set when you set up the storage system.

Storage-system management control properties

The storage-system management control properties are:

- ◆ Management control
- ◆ Management control status
 - Enable control
 - Disable control

Management control

Management control enables or disables management control for the selected storage system. Management control is available for pre-FC4700 shared storage systems only. The first person to enable management control for a storage system must define a management login password. Thereafter, anyone who disables management control for the storage system or changes management control status for the storage system must enter this password.

By default, management control is disabled, so any server connected to a storage system can send configuration commands to the storage system. If you enable management control, none of the servers connected to the storage system can send configuration commands to it because its storage-system control status is disabled. Use the enable control property to permit one or more servers to send configuration commands to the storage system.

Management control status - enable control, disable control

Enable control enables management control for the servers that you select. Enabling management control lets the selected servers send configuration commands to the storage system. If management control is enabled for a server, it is enabled for all server initiators (HBA ports) connected to the storage system. You must enter the management login password before enabling management control for the servers.

Disable control disables management control for the server that you select. Disabling management control prevents the selected servers from sending commands to the storage system. If management control is disabled for a server, it is disabled for all server initiators (HBA ports) connected to the storage system. You do not need to enter the management login password before disabling management control for the servers.

All servers can send certain LUN configuration commands to the storage system even when management control is disabled for them. These commands set the user-defined properties on the **General**, **Cache**, and **Prefetch** tabs in the **LUN Properties** dialog box, which are the properties listed below.

Tab	User-defined properties
General	Rebuild Priority, Verify Priority, Auto Assignment Enabled, Default SP Owner
Cache	Read Cache Enabled, Write Cache Enabled
Prefetch	All properties

Enabling management control for a pre-FC4700 shared storage system

Before you can enable or disable management control for the storage system from a server, you must enable management control for the storage system, which involves setting the management login password for the storage system.

If management control is *not* enabled for a storage system, any server connected to the storage system can send configuration commands to the storage system.

You need the management login password to enable management control for hosts or to disable management control for the storage system. If no one can remember the current password, then you must connect a management station to the serial port on a storage-system SP to change the password without having to enter the current password.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system on which you want to enable configuration access control, and then click **Properties**.
2. In the **Management Control** tab, select the **Management Control** check box, and then click **OK**.

3. Enter a new password and then confirm that password.

Management control status for all hosts connected to this storage system is set to disabled, and no hosts are able to issue configuration commands to the storage system. You must enable management control for specific hosts.

What next? Continue to the next section to enable management control for servers.

Enabling and disabling management control for servers

The first time that management control is enabled for the storage system, management control is disabled for all servers connected to it. Thereafter, if storage-system management control is disabled and then re-enabled, the management control for a server is the same as it was when storage-system management control was disabled.

When you enable management control for a server, it is enabled for each server HBA port (initiator) connected to the storage system.

To enable management control for servers

IMPORTANT: Before you can enable management control for a server to the storage system, you must enable management control for the storage system (refer to page 7-4).

1. Open the **Management Control** tab of the **Storage System Properties** dialog box. If it is *not* open, refer to "Enabling management control for a pre-FC4700 shared storage system" on page 7-4 for instructions on how to do so.
2. Under **Management Control Status**, select the servers that should have management control for the storage system and click **Enable Control**.
3. In the **Enable Management Login** dialog box, enter the management login password.

The **Management Control Status** for the selected host changes from **Disabled** to **Enabled**.

To disable management control for servers

1. Select the **Management Control** tab of the **Storage System Properties** dialog box for the storage system.
2. Under **Management Control Status**, select the servers that should not have configuration access to the storage system and click **Disable Control**.

The **Management Control Access** for the selected hosts changes from **Enabled** to **Disabled**.

The servers with access enabled can now send configuration commands to the storage system. The servers with access disabled cannot send configuration commands to the storage system.

What next?

Continue to the next section, "Setting storage-system general configuration properties".

Setting storage-system general configuration properties

This section describes the general configuration properties for a storage system and explains how to set them.

General configuration properties

The general configuration properties are:

- ◆ SP A statistics logging
- ◆ SP B statistics logging
- ◆ Assigning a custom to a storage system

SP A and SP B statistics logging

SP A statistics logging enables or disables logging of statistics by SP A, and SP B statistics logging enables or disables logging of statistics by SP B.

Each SP maintains a log of statistics for the LUNs, disks, and storage-system caching. You can turn this log on or off.

If Navisphere Analyzer is installed, enabling statistics logging also turns Analyzer logging on.

Storage-system descriptions

The default storage-system description has the following format:

storage_system_serial#[type]

where

storage_system_serial# is the unique serial number of enclosure 0 in a CX-Series or FC-Series storage system or the chassis in a C-series storage system.

type is the storage-system type; for example, 95-2694-261 [FC4700]

Depending on the status of the storage system, the storage system type may be replaced with *Inaccessible* or *Unsupported*; for example, 95-2694-261 [Inaccessible]

If you want to replace the storage-system serial number (the default value) with a custom name, go to the next section, "Setting the general configuration properties".

Setting the general configuration properties

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose configuration properties you want to set or change, click **Properties**, and then select the **General** tab.
2. You can set or change the following general configuration properties:
 - Storage-system name, serial number, and Base UUID.

Changing the name does not affect the agent configuration file.

If different storage systems have the same serial number, Navisphere Manager may exhibit strange behavior. If a storage system has no serial number, Navisphere Manager 6.X cannot manage the storage system. Make sure that each storage system has a unique serial number.

- SP A and SP B logging statistics

Refer to the online help for descriptions of these properties.

3. You can view the **Management Connections** list to verify the host connections to the storage system.

What next?

If you need to allocate memory for caching or for binding RAID 3 LUNs continue to the next section, "Setting storage-system memory properties".

Setting storage-system memory properties

Setting memory properties consists of assigning memory to memory partitions. You must assign memory to the appropriate partitions to do the following:

- ◆ Use read or write caching
- ◆ Bind RAID 3 LUNs

This section describes:

- ◆ Memory requirements for read and write caching
- ◆ Memory requirements for RAID 3 LUNs
- ◆ Effects of SP memory architecture on memory assignment
- ◆ How to assign memory to partitions



CAUTION

Before you bind a RAID 3 LUN do the following:

For pre-FC4700 storage systems

Assign at least 2 Mbytes per RAID 3 LUN to the RAID 3 memory partition. If this partition does not have adequate memory for the LUN, you will not be able to bind it. Changing the size of the RAID 3 memory partition reboots the storage system. Rebooting restarts the SPs in the storage system, which terminates all outstanding I/O to the storage system.

For CX-Series or FC4700-Series storage systems that support RAID 3 LUNs

Allocating the RAID 3 memory partition size is not required for RAID 3 Raid Groups and LUNs. (The RAID 3 memory partition appears dimmed and is unavailable.) If there will be a large amount of sequential read access to this RAID 3 LUN, you may want to enable read caching with prefetching for the LUN.

Assigning memory to partitions

Use the **Memory** tab on the storage-system **Properties** dialog box to assign storage-system memory on each SP to these partitions:

- ◆ Read cache
- ◆ Write cache
- ◆ RAID 3 partitions

These partitions have a default size of 0.



CAUTION

For pre-FC4700 storage systems, changing the RAID 3 partition size causes the application to reboot the storage system. This terminates all outstanding I/O to the storage system.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose memory you want to assign, and then click **Properties**.
2. Select the **Memory** tab, and under **User Customizable Partitions**, enter a new value to adjust the size of each memory partition that you want to change. You can change the SP A and SP B read cache memory, and the write cache memory.

When you do this, Manager reassigns memory in one of two ways:

- *From* free memory to a partition whose size you are increasing
- *To* free memory from a partition whose size you are decreasing

The pie charts reflect the changes in memory assignment.

As a general guideline, we recommend that you make the write-cache partition approximately twice the size of the read-cache partition on each SP. For example, if total memory for each SP is 256, you can assign 150 to the write-cache partition and 75 to the read-cache partition on each SP. For systems with modest prefetch requirements, 50 MB to 100 MB of read cache per SP is sufficient.

What Next? Your next action depends on whether you assigned memory to the read-cache or write-cache memory partitions.

Memory assigned to cache partitions - Continue to the next section, "Setting storage-system cache properties".

Memory not assigned to cache partitions - Go directly to Chapter 8, which describes how to create LUNs.

Setting storage-system cache properties

This section describes:

- ◆ Hardware requirements for storage-system caching
- ◆ Storage-system cache properties
- ◆ Setting the cache properties for a storage system

All storage systems support read caching. A storage system supports write caching only if it has the required hardware, which varies with the storage-system type as shown in the following table.

Table 7-1 Hardware requirements for write caching

Icon	CX-series	FC4400/4500, FC4700, FC5600/5700	FC5200/5300	C1900, C2x000, C3x00	C1000
Disks	0-0 through 0-4	0-0 through 0-8	0-0 through 0-4	A0, B0, C0, D0, E0	A0 through A4
SPs	Two	Two		Two, with at least 8 Mbytes memory	
Power supplies	Two in SPE or DPE2, and each DAE2	Two in DPE and each DAE		Two	
LCCs	Two in each DAE2	Two in DPE and each DAE		Not applicable	
Backup power	Fully charged SPS	Fully charged SPS		Fully charged BBU	

Storage-system cache properties

Page size Page size sets the number of Kbytes stored in one cache page. The storage processors (SPs) manage the read and write caches by pages instead of sectors. The larger the page size, the more continuous sectors the cache stores in a single page. The default page size is 2 Kbytes.

As a general guideline, we recommend the following page sizes:

- ◆ For general file server applications: 8 Kbytes
- ◆ For database applications: 2 or 4 Kbytes

Low watermark, high watermark, enable watermark processing

The SPs use high and low watermarks to determine when to flush their write caches. When an SP flushes its write cache, it writes its dirty pages to disk. A dirty page is a write-cache page with modified data that has not yet been written to disk.

You can make the following selections regarding watermark processing:

Selection	Function	Action	Remarks
Low watermark	Sets the low watermark	In the Low Watermark box, select a value for the low watermark	Determines when flushing stops
High watermark	Sets the high watermark	In the High Watermark box, select a value for the high watermark	Determines when flushing starts
Enable watermark processing	Enables or disables watermark processing	Check the Enable Watermark Processing box to enable watermark processing	If you do not enable watermark processing, the system sets both watermarks to 100

Following are further details about high and low watermarks:

Water-mark type	Definition	Impact of high or low values	Default value
High	Percentage of dirty pages in the write cache which, when reached, causes the SPs to begin flushing their write caches	A low value for the high watermark causes the SPs to begin flushing their write caches sooner than a high value	FC4700-60% CX-Series-80%
Low	Percentage of dirty pages in the write cache which, when reached, causes the SPs to stop flushing their write caches	A high value for the low watermark causes the SPs to stop flushing their write caches sooner than a low value	FC4700-40% CX-Series-60%

Mirrored write cache

Mirrored write cache sets the write-cache type for both SPs in a storage system to either mirrored or non-mirrored. If the SP with a non-mirrored write cache fails, all write-cache data not saved to disk is lost. To provide data integrity, the write-cache type is mirrored for the SPs in most types of storage systems.

SP A and SP B read caching

You can make the following selections for SP A and SP B read caching:

Type of read caching	Function	Method
SP A	Enables or disables storage-system read caching for SP A	Enables or disables the read cache on SP A
SP B	Enables or disables storage-system read caching for SP B	Enables or disables the read cache on SP B

The read cache on one SP is independent of the read cache on the other SP. On powerup, a storage system automatically enables the storage-system read caching for an SP if the SP's read-cache size is non-zero.

You can enable or disable storage-system read caching for an SP without affecting the information on the LUNs owned by the SP. You must enable the storage-system read caching for an SP before any of the SP's LUNs with read caching enabled can use read caching.

Some other operations, such as setting most of the LUN caching properties, require that the SP A or SP B read cache be disabled. If the cache is enabled when you perform any of these operations, Manager automatically disables the appropriate cache for you and re-enables it after the operation is complete.

Write caching

Write caching enables or disables storage-system write caching by enabling or disabling the write cache on each SP. The write cache on one SP mirrors the write cache on the other SP. As a result, both write caches are always the same size and are always either both enabled or both disabled. On powerup, a storage system automatically enables the write cache on each SP if the write-cache size is non-zero.

You can enable or disable storage-system write caching without affecting the information on the LUNs owned by the SP. You must enable the storage-system write caching before any LUNs with write caching enabled can use write caching.

Some other operations, such as setting most of the LUN caching properties, require that the write cache be disabled. If the write cache is enabled when you perform any of these operations, Manager automatically disables the write cache for you and re-enables it after the operation is completed.

Navisphere Manager lets you determine the availability of storage-system write caching when a single drive in the cache vault fails. When you select the **HA Cache Vault** check box (this is the default) in the **Storage System Properties - Cache** tab, write caching is disabled if a single vault disk fails. When you clear the check box, write caching is not disabled if a single disk fails.

If you do not disable write caching when a single cache vault disk fails, the data is at risk if another cache vault disk should fail.

Setting the cache properties

Use the **Cache** tab on the **Storage System Properties** dialog box to set the storage-system cache properties.

The minimum write cache and read cache partition size is 1 Mbyte for pre-FC4700 storage systems, and 3 Mbytes for CX-Series and FC4700-Series storage systems.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose cache properties you want to change, click **Properties**, and then select the **Cache** tab.
2. Under **Configuration**, you can do any of the following:
 - Set the page size
 - Enable or disable watermarks and change the values for low and high watermark.
 - Enable or disable SP A and SP B read cache.
 - Enable or disable mirrored write cache.
 - Enable or disable storage-system write caching.
 - Enable or disable storage-system write caching if a cache vault disk fails.

If storage-system read or write caching is enabled when you change the value of a cache property, then read or write caching is disabled before the change is made, and then re-enabled when the change is complete.

If an SP is removed, you cannot use Navisphere Manager to set the write cache. Use Navisphere CLI.

Setting the fair access property (pre-FC4700 only)

This section:

- ◆ Describes fair access to storage-system resources.
- ◆ Explains how to set the enforce fair access property for a pre-FC4700 storage system.

Fair access to the storage-system resources

By default, for pre-FC4700 storage systems, the storage system processes I/O requests from servers on a first-come, first-served basis. With multiple servers contending for the use of a storage system, a server with disproportionate processing demands might monopolize storage-system resources. In addition, operating systems, such as Windows, use scheduling policies that slow down shared storage-system access.

To provide each server connected to shared storage systems with a fair amount of storage-system resources, shared storage systems have an optional fairness algorithm. This algorithm tries to manage the I/Os accepted by the storage system so that servers accessing different LUNs with similar data access patterns will get similar I/O throughput. Some data access patterns, however, do not work well with the algorithm.

We strongly recommend that you try using fair access, especially if Windows servers are accessing the storage system. Should I/O performance be unsatisfactory, you can turn off fair access and return to the first-come, first-served algorithm.

Enabling fair access to a storage system

If the selected storage system does not support the fair access property, **Enforce Fair Access** will appear dimmed and be unavailable.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate the storage-system icon for which you want to enable fair access, and then click **Properties**.
2. In the **Storage Access** tab, select the **Enforce Fair Access** check box.

What next? **For a CX-Series or FC4700-Series storage system** - If you have not set the IP address for the SPs and ALPAs, go to the next section, *"Setting the SP network and ALPA properties (CX-series or FC4700-series only)"*. If you have set them and you plan to change the battery and use caching, go to the section "Setting the battery test time" on page 7-20. Otherwise, go directly to Chapter 8 to create LUNs on the storage system.

For a pre-FC4700 storage system - If you plan to change the battery and use caching, go to the section "Setting the battery test time" on page 7-20; otherwise, go directly to Chapter 8 to create LUNs on the storage system.

Setting the SP network and ALPA properties (CX-series or FC4700-series only)

The SP network properties establish the network name and address for each SP; the ALPA (Arbitrated Loop Physical Address) properties establish the SCSI ID for each SP port. The settings of these properties must be correct; if any of the properties is incorrect, Manager cannot communicate with the SP and its LUNs.

Setting the SP network properties

Important Before changing a storage system's IP address, EMC strongly recommends that you remove the storage system from its domain. Once you assign the new IP address, you can safely return the storage system to the domain.

The SP Network properties include:

- ◆ SP hostname (used when you select the SP for management)
- ◆ SP IP address, subnet address, and network mask (required to let the management station use the Internet connection to communicate with the SP)

The **SP Properties - Network** tab lets you manage the SP Internet connection by changing the SP network name and address.

The network properties are initially set by EMC service personnel to work at your site. Do not change any value unless you are moving the SP to another LAN or subnet. If you change any value, after you click **OK** or **Apply**, the SP will restart and use the new value.

To Set the SP network properties

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the SP whose properties you want to change, click **Properties**, and then select the **Network** tab.
2. Specify the new network name and address settings you want, then click **OK** or **Apply**.

The SP restarts using the new values specified.

What next?

The SP network properties are independent of other SP properties; there is no related setting you need to change next. Depending on your reason for changing this SP's network properties, you may want to change one or more network properties of the other SP in this storage system.

Setting the SP ALPA (Arbitrated Loop Physical Address) properties

The **SP ALPA Properties** dialog box lets you change the SCSI ID (ALPA address) of each SP port.

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

If you change any value, after you click **OK** or **Apply**, the SP will restart and use the new values.

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

The software will not let you select a SCSI ID out of range (0-126) or a duplicate ID on a storage system.

To set the SCSI ID associated with an SP's port

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the SP whose properties you want to change, click **Properties**, and then select the **ALPA** tab.
2. Click **Change** to open the **Change ALPA Setting** dialog box.
3. Specify a new SCSI ID for one SP port, and then click **OK**.

The SP restarts using the new values specified.

What next?

The SP port ALPA addresses (SCSI IDs) are independent of other SP properties; there is no related setting you need to change next. Depending on your reason for changing this SP's port SCSI IDs, you may want to change the IDs of the other SP in this storage system.

If you are using caching and want to change the battery test time, continue to the section, "Setting the battery test time"; otherwise, go directly to Chapter ""8 to create LUNs on the storage system.

Setting the SP fibre speed property

The **Fibre Speed** tab is visible only for CX-Series or FC4700-Series storage systems that support 2-Gbit fibre connections.

Use the **Fibre Speed** tab on the **SP Properties** dialog box to change the fibre speed for each SP port.

Setting the SP port speed

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the SP component for which you want to set the SP port speed, click **Properties**, and then select the **Fibre Speed** tab.
2. Under **Current Settings**, select the speed for each port before saving your changes and closing the dialog box.

Adding privileged users to the SP agent configuration file

To add privileged users to the SP Agent configuration file, do the following:

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icon for the storage system in which the SP resides.
3. Right-click the icon for the SP (A or B) that you want to edit the configuration file, and click the **Agent** tab.
4. Under **Privileged Users**, click a blank line, and then enter the name of the privileged user.
5. Click **OK** to save your changes and close the dialog box.

Setting the battery test time

Refer to the online help for detailed instructions on setting battery test time.

Each week, the SP runs a battery self-test to ensure that the monitoring circuitry is working in each SPS in an FC-series storage system or in the BBU of a C-series storage system.

While the test runs, storage-system write caching is disabled, but communication with the server continues. I/O performance may decrease during the test. When the test is finished, storage-system write caching is re-enabled automatically.

The factory default setting has the battery test start at 1:00 a.m. on Sunday. You can change this setting using the procedure that follows.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for SPS A or SPS B or the BBU whose battery test time you want to change and click **Set Test Time**.
2. In the **Battery Test Time** dialog box, fill in the day and time on which you want the test to run.

For legacy storage systems, or storage systems with pre-6.5.0 Agents, convert the local time to GMT time and enter the GMT value.

Registering HBAs and setting HBA initiator record values

Navisphere Manager lets you register individual HBAs connected to a specific storage system with that storage system, or you can register all HBAs from a specific server to all connected storage systems.

Registering HBAs with a specific storage system

The **Create Initiator Record** dialog box allows you to permanently register a new or unregistered HBA with the SP port to which it is connected. When you register an HBA, the record for each HBA initiator is added to the storage system's server initiator table, and the Agent then recognizes the HBA. When all HBAs belonging to a server are registered, the server can be connected to and disconnected from Storage Groups.

For supported storage systems, different server failover software (for example, ATF and PowerPath® software) requires different values for **Initiator Type**, **ArrayCommPath** and **Failover Mode**. If you are using failover software, go to Chapter 5, "Setting Storage System Failover Properties".

If you are *not* using failover software, set the storage-system failover properties as follows:

- ◆ Unless instructed otherwise, set **Initiator Type** to **CLARiiON Open**.
- ◆ Enable **ArrayCommPath**.
- ◆ Set **Failover Mode** to 0.



CAUTION

If you enter incorrect failover values the storage system could become unmanageable or unreachable by the server, and the failover software could operate incorrectly.

Registering an existing unregistered HBA

To register an existing unregistered HBA, do the following:

1. In the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose server initiator you want to register, and click **Connectivity Status** to open the **Connectivity Status** dialog box.
2. In **Initiator Records**, select the HBA that you want to register, and then click **Register** to open the **Register Initiator Record** dialog box.

The selected HBA must have a **No** value in the **Registered** column to enable **Register**.

For a description of the properties in the dialog box, click **Help**.

3. Verify that the **HBA WWN** and the **SP Port ID** are correct.
4. Unless instructed otherwise, set **Initiator Type** to **CLARiiON Open**.
5. Enable or disable **Array CommPath**.
6. In **Failover Mode**, select the correct failover value.
7. Unless the storage system is connected to a Solaris server, set **Unit Serial Number** to **Array**.
8. Enter the **Vendor** and **Model** for the initiator.
9. Under **This HBA Belongs to**, click **Existing Host**, and then from the list, select the host in which the HBA resides.

When registering Tru64 HBAs to the storage system, the Tru64 servername may not display in the **This HBA belongs to** drop-down menu.

10. Save your changes and close the dialog box to return to the **Connectivity Status** dialog box.

Registering a new unregistered HBA

To register a new unregistered HBA, do the following:

1. From the **Connectivity Status** dialog box, click **New** to open the **Create Initiator Record** dialog box.

For a description of the properties in the dialog box, click **Help**.

2. In the **Create Initiator Record** dialog box, enter the **HBA WWN** (World Wide Name) and **SP Port ID** (0 or 1) for the new HBA.
3. Unless instructed otherwise, set **Initiator Type** to **CLARiiON Open**.
4. Enable or disable **Array CommPath**.
5. **In Failover Mode**, select the correct failover value.
6. Under **This HBA Belongs to**, click **New Host**, and then enter a valid **Host Name** and enter the **IP Address** for the host in which the new HBA resides.
7. Save your changes and close the dialog box to return to the **Connectivity Status** dialog box.

Registering HBAs from a specific host

The **Connect Hosts** dialog box from the **Tools** menu scans the host system and displays information about the devices to which the selected host is able to communicate. Use the **Connect Host** dialog box to register all HBAs in the selected host with all connected storage systems, and to diagnose server related problems.

In order to use the tool, the following must be true:

- ◆ The Host Agent configuration file includes the IP address of the managing SP (the SP to which you are pointing your browser). Add the IP address using the **system@*a.b.c.d*** format where *a.b.c.d* is the IP address for the managing SP.
- ◆ The host is connected to the storage systems.
- ◆ The Host Agent, connected to the storage systems, is running.
- ◆ The host is managed, either automatically or using a portal.
- ◆ The managing storage system is running version 6.6 or higher of the Management Server Software.
- ◆ The host is running version 6.6 or higher of the Host Agent.

Scanning the host system for paths and volumes

Use this method to view all available device paths and addressable volumes for the specified host.

1. In the **Tools** menu, click **Connect Host** to open the **Connect Host** dialog box.
2. Enter the IP address for the host you want to register, scan or diagnose server problems, and click **Go**.

Manager displays the data path and CLARiiON volume information for the host.

For more details about the displayed information, see the online help for the **Connect Host** dialog box.

Registering hosts

Registers all HBAs in the host with the connected storage systems.

1. In the **Tools** menu, click **Connect Host** to open the **Connect Host** dialog box.
2. Enter the IP address for the host you want to register and click **Go**.

Manager displays the data path and CLARiiON volume information for the host.

3. Click **Register** to register the host with all connected storage systems.

Updating data path information

Refreshes the data path information.

Scan does not update any volume information. To do this, click **Go**, or restart the Host Agent.

In the open **Connect Host dialog** box, click **Scan**.

Manager updates the data path information for the host.

To deregister an HBA initiator

Use the **Deregister** button to remove ghost hosts from the initiator table. Ghost hosts are hosts that were disconnected from a storage system, but the storage system still sees the hosts as being registered.

You can identify a ghost host record as follows:

- ◆ The icon preceding the initiator record is dimmed
 - ◆ The server name is unrecognizable
 - ◆ The IP address and initiator WWN belong to hosts that you know have been disconnected
 - ◆ The initiators are still registered but not *logged* in.
1. In the **Enterprise Storage** dialog box, navigate to the icon for the storage system whose host initiator you want to deregister, and click **Connectivity Status** to open the **Connectivity Status** dialog box.
 2. In **Initiator Records**, select the record for the HBA that you want to deregister, and then click **Deregister**.
 3. In the confirmation dialog box, click **Yes** to continue or **No** to cancel the deregister operation.

Manager removes the initiator record from the dialog box.

What next? Continue to Chapter 8 to create LUNs on the storage system.

Creating LUNs and RAID Groups

You can create LUNs on any storage system with SPs, that is, any storage system except an FC5000 series storage system (JBOD configuration). You can create RAID Groups on any storage system that supports the RAID Group feature.

For security reasons, we recommend that you place all storage systems in a domain before you configure them.

To configure the storage system, you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes the following:

- ◆ About LUNs.....8-2
- ◆ About RAID groups..... 8-11
- ◆ Creating RAID groups.....8-14
- ◆ Creating LUNs on RAID groups 8-16
- ◆ Binding a new LUN with the exact size of an existing LUN.....8-20
- ◆ Creating LUNs on a non-RAID group storage system.....8-21

About LUNs

This section describes the following topics:

- ◆ LUNs, the RAID types for LUNs, and the LUN properties.
- ◆ RAID Groups and RAID Group properties.

LUNs

A logical unit (LUN) is a grouping of one or more disks or disk partitions into one span of disk storage space. A LUN looks like an individual disk to the server's operating system. It has a RAID type and properties that define it.

You can have Manager create standard LUNs using the disks and default property values that it selects, or you can create your own custom LUNs with the disks and property values that you select. In a storage system that supports RAID Groups, you create LUNs on RAID Groups; therefore, you need to create a RAID Group before you create a LUN.

If you need a LUN that exceeds the maximum number of disks for a RAID type, or you need to expand the user capacity of an existing LUN, use the metaLUN feature of Navisphere Manager (see Chapter 10, "Expanding LUN Capacity With MetaLUNs")

RAID types

The RAID type of a LUN determines the type of redundancy, and therefore, the data integrity provided by the LUN.

The following RAID types are available:

RAID 5 - A distributed parity array, which provides data integrity using parity information that is stored on each disk in the LUN. This RAID type is best suited for multiple applications that transfer different amounts of data in most I/O operations.

RAID 3 - A single-disk parity array, which provides data integrity using parity information that is stored on one disk in the LUN. This RAID type is best suited for single-task applications, such as video storage, that transfer large amounts of data in most I/O operations.

RAID 1 - A mirrored array, which provides data integrity by mirroring (copying) its data onto another disk in the LUN. This RAID type provides the greatest data integrity at the greatest cost in disk space, and is well suited for an operating system disk.

RAID 1/0 - A mirrored individual access array without parity, which provides the same individual access features as the RAID 5 type, but with the highest data integrity. This RAID type is well suited to the same applications as the RAID 5 type, but where data integrity is more important than the cost of disk space.

RAID 0 - An individual access array without parity, which provides the same access features as the RAID 5 type, but does not have parity information. As a result, if any failure (including an unrecoverable read error) occurs on a disk in the LUN, the information on the LUN is lost.



CAUTION

Since the RAID 0 type provides no protection for your data, EMC does not recommend using this RAID type unless you have some way of protecting your data, such as software mirroring.

Disk - An individual disk type, which functions just like a standard single disk, and, as such, does not have the data integrity provided by parity or mirrored data. This RAID type is well suited for temporary directories that are not critically important.

Hot Spare - A single global spare disk, which serves as a temporary replacement for a failed disk in a RAID 5, 3, 1, or 1/0 LUN. Data from the failed disk is reconstructed automatically on the hot spare. It is reconstructed from the parity data or mirrored data on the working disks in the LUN; therefore, the data on the LUN is always accessible. A hot spare LUN cannot belong to a Storage Group.

Number of disks in a LUN

The RAID type of a LUN determines the number of disks that you can select for the LUN, shown as shown in Table 8-1.

If you need a LUN that exceeds the maximum number of disks for a RAID type, or you need to expand the user capacity of an existing LUN, use the metaLUN feature of Navisphere Manager (see Chapter 10, "Expanding LUN Capacity With MetaLUNs").

Table 8-1 Number of disks you can use in RAID types

RAID type	Number of disks you can use
RAID 5	3 - 16
RAID 3	5 or 9 (CX-Series or FC-Series)
RAID 1/0	2, 4, 6, 8, 10, 12, 14, 16
RAID 1	2
RAID 0	3 - 16
Disk	1
Hot Spare	1

Disks that cannot be hot spares

Table 8-2 lists the disks in each type of storage system that you cannot use as hot spares.

Table 8-2 Disks that cannot be hot spares

Series	Disk IDs
AX-Series	0:0 through 0:3
CX700,CX600, CX500,CX500i, CX400,CX300, CX200,CX200LC	0:0 through 0:4
FC4700-Series FC4400/4500, FC5500/5600	0:0 through 0:8

Table 8-2 Disks that cannot be hot spares (*continued*)

Series	Disk IDs
FC5200/5300	0:0 through 0:4
C3x00	A0, B0, C0, D0, E0, A3, A4
C2x00	A0, B0, C0, D0, E0, A3
C1900	A0, B0, C0, D0, E0, A1
C1000	A0, A1, A2, A3, A4, B0

Do not use an ATA (Advanced Technology Attachment) disk as a spare for a Fibre-Channel-based LUN, and do not use a Fibre Channel disk as a spare for an ATA-based LUN.

LUN properties

The LUN properties determine the individual characteristics of a LUN. You set LUN properties when you bind the LUN and you can change certain LUN properties after the LUN is bound. The LUN properties are as follows:

- ◆ LUN ID (assigned at creation; cannot be changed)
- ◆ LUN Size (RAID Group storage systems only)
- ◆ Rebuild Priority
- ◆ Verify Priority
- ◆ Default Owner
- ◆ Enable Read Cache
- ◆ Enable Write Cache
- ◆ Enable Auto Assign
- ◆ Number of LUNs to bind
- ◆ Alignment offset (available only for CX-Series and FC4700-Series storage systems with a specific FLARE software revision. Refer to the Manager release notes)

LUN properties are not available for the Hot Spare RAID type because it is simply a replacement disk for a failed disk in a LUN.

Rebuild priority

The rebuild priority is the relative importance of reconstructing data on either a hot spare or a new disk that replaces a failed disk in a LUN. It determines the amount of resources the SP devotes to rebuilding instead of to normal I/O activity. Table 8-3 lists and describes the rebuild time associated with each rebuild value.

Table 8-3 Valid rebuild priorities

Value	Target rebuild time in hours
ASAP	0 (as quickly as possible) This is default.
HIGH	6
MEDIUM	12
LOW	18

The rebuild priorities correspond to the target times listed above. The storage system attempts to rebuild the LUN in the target time or less. The actual time to rebuild the LUN depends on the I/O workload, the LUN size, and the LUN RAID type.

For a RAID Group with multiple LUNs, the highest priority specified for any LUN on the group is used for all LUNs on the group. For example, if the rebuild priority is High for some LUNs on a group and Low for the other LUNs on the group, all LUNs on the group will be rebuilt at High priority.

You set the rebuild priority for a LUN when you bind it, and you can change it after the LUN is bound without affecting the data on the LUN.

Verify priority

The verify priority is the relative importance of checking parity sectors in a LUN. If an SP detects parity inconsistencies, it starts a background process to check all the parity sectors in the LUN. Such inconsistencies can occur after an SP fails and the other SP takes over the LUN. The priority determines the amount of resources the SP devotes to checking parity instead of to normal I/O activity.

Valid verify priorities are ASAP (as soon as possible – the default), HIGH, MEDIUM, and LOW. A verify operation with an ASAP or HIGH priority checks parity faster than one with a MEDIUM or LOW priority, but may degrade storage-system performance.

You set the verify priority for a LUN when you bind it, and you can change it after the LUN is bound, without affecting the data on the LUN.

Default owner The default owner is the SP that assumes ownership of the LUN when the storage system is powered up. If the storage system has two SPs, you can choose to bind some LUNs using one SP as the default owner and others using the other SP as the default owner. The primary route to a LUN is the route through the SP that is its default owner, and the secondary route is through the other SP. LUNs that are not currently owned by an SP are unowned. A hot spare that is not in use is an unowned LUN.

Valid default owner values are SP A, SP B, and Auto. Auto attempts to divide the LUNs equally between SP A and SP B. The default value is SP A for a storage system with one SP, and Auto for a storage system with two SPs.

You set the default owner for a LUN when you bind it and you can change it after the LUN is bound, without affecting the data on the LUN.

Enable read cache Enable read cache enables (default) or disables read caching for a LUN. For a LUN with read caching enabled to actually use read caching, the read cache on the SP that owns the LUN must also be enabled. If the read cache for the SP owning the LUN is enabled, then the memory assigned to that read cache is shared by all LUNs that are owned by that SP and have read caching enabled.

Generally, you should enable read caching for every RAID type that supports it. If you want faster read performance on some LUNs than on others, you may want to disable read caching for the lower priority LUNs.

You enable or disable read caching for a LUN when you bind it. You can also enable or disable read caching after the LUN is bound without affecting its data.

Enable write cache Enable write cache enables (default) or disables write caching for a LUN. For a LUN with write caching enabled to actually use write caching, the write cache for the storage system must also be enabled. If the storage-system write cache is enabled, then the memory assigned to the write cache is shared by all LUNs that have write caching enabled.

Generally, you should enable write caching for every RAID type (especially for a RAID 5 or RAID 1/0 LUN) that supports it. If you want faster write performance on some LUNs than on others, you may want to disable write caching for the lower priority LUNs.

You enable or disable write caching for a LUN when you bind it. You can also enable or disable write caching after the LUN is bound, without affecting its data.

Enable auto assign

Enable auto assign enables or disables (default) auto assignment for a LUN. Auto assignment controls the ownership of the LUN when an SP fails in a storage system with two SPs.

You enable or disable auto assignment for a LUN when you bind the LUN. You can also enable or disable it after the LUN is bound, without affecting the data on it.

With auto assignment enabled, if the SP that owns the LUN fails and the server tries to access that LUN through the second SP, the second SP assumes ownership of the LUN so the access can occur. The second SP continues to own the LUN until the failed SP is replaced and the storage system is powered up. Then, ownership of the LUN returns to its default owner.

If auto assign is disabled in the previous situation, the other SP does not assume ownership of the LUN, so the access to the LUN does not occur.

If you are running failover software on a server connected to the LUNs in a storage system, you must *disable* auto assignment for all LUNs that you want the software to fail over when an SP fails. In this situation, the failover software, not auto assign, controls ownership of the LUN in a storage system with two SPs.

Alignment offset

Alignment offset is available only for CX-Series and FC4700-Series storage systems with a specific revision of FLARE software. Refer to the Manager release notes.

Alignment offset automatically aligns the logical sector number of the metaLUN to a stripe boundary on the base LUN. The value to enter is the number of sectors by which user data is offset due to file system or partition metadata. Correctly setting the alignment offset value should improve performance. Problems can arise when a server operating system records private information at the start of a LUN. This can interfere with the RAID stripe alignment so that when data I/O crosses the RAID stripe boundary the storage system performance is degraded. Table 8-4 lists the LUN properties available for different RAID types, and Table 8-5 lists the default properties for the RAID types.

No initial verify

When cleared (this is the default), the software performs an initial background verify on the newly bound LUN, and eliminates any latent soft media errors. When enabled, the software does not perform the initial background verify.

Do not send data to the LUN until the background verify operation is complete.

Table 8-4 LUN properties available for different RAID types

Property	RAID 5	RAID 3	RAID 1/0	RAID 1	RAID 0	Disk
LUN ID	Yes	Yes	Yes	Yes	Yes	Yes
LUN size (RAID Group LUNs only)	Yes	Yes*	Yes	Yes	Yes	No
Rebuild priority	Yes	Yes	Yes	Yes	No	No
Verify priority	Yes	Yes	Yes	Yes	No	No
Default owner	Yes	Yes	Yes	Yes	Yes	Yes
Enable read cache	Yes	Yes*	Yes	Yes	Yes	Yes
Enable write cache	Yes	Yes*	Yes	Yes	Yes	Yes
Enable auto assign	Yes	Yes	Yes	Yes	Yes	Yes
Alignment offset	Yes	Yes	Yes	No	Yes	No
No initial verify	Yes	Yes	Yes	Yes	Yes	Yes
* Supported only on storage systems running release 13 or later FLARE software						

Table 8-5 Default LUN property Values for different RAID types

Property	RAID 5	RAID 3	RAID 1/0	RAID 1	RAID 0	Disk
LUN ID	Smallest available number					
LUN size	GB is the default unit of measure. The default value is 1 for GB; 100 for MB; maximum number of blocks for a LUN of the selected RAID type for Block Count.					
Rebuild priority	ASAP	Low	Low	Low	N/A	N/A
Verify priority	ASAP	Low	Low	Low	N/A	N/A
Default owner	SP A for one SP; Auto for two SPs. Auto distributes the LUNs as equally as possible between the two SPs.					
Enable read cache	Selected	Selected	Selected	Selected	Selected	Selected
Enable write cache	Selected	Selected	Selected	Selected	Selected	Selected
Enable auto assign state	Cleared	Cleared	Cleared	Cleared	Cleared	Cleared
Alignment offset	0	0	0	N/A	0	N/A

Table 8-6 Recommendations for creating LUNs with different RAID types

RAID type	Recommendations
RAID 5	Select 5 to 10 disks to achieve a balance between usable capacity and rebuild time. Rebuild time increases with the number of disks selected.
RAID 3	Select 5 drives for quicker rebuild times or 9 drives for more efficient use of space.
RAID 1/0	Disks are paired into mirrored images in the order in which you select them. The first and second disks you select are a pair of mirrored images; the third and fourth disks you select are another pair of mirrored images; and so on. For highest data integrity, the first disk you select in each pair should be on a different disk bus than the second disk you select.
RAID 0	None
RAID 1	None
Disk	None
Hot Spare	None

About RAID groups

A RAID Group is a set of disks on which you bind one or more LUNs. Each LUN you bind on a RAID Group is distributed equally across the disks in the Group.

You do not define the RAID type of the RAID Group when you create it. The RAID Group supports the RAID type of the first LUN you bind on it. Any other LUNs that you bind on it have the same RAID type. The number of disks you select for this RAID Group determines the RAID types this RAID Group will support. For example, if you select four disks for your new RAID Group, this RAID Group will support the RAID 0, RAID 5, and RAID 1/0 RAID types.

RAID group expansion

You can expand a RAID Group by adding one or more disks to it. Expanding a RAID Group does *not* automatically increase the user capacity of already bound LUNs. Instead, it distributes the capacity of the LUNs equally across all the disks in the RAID Group, freeing space for additional LUNs.

If you expand a RAID Group that has only *one* bound LUN with a user capacity *equal* to the user capacity of the RAID Group, you can choose to have the user capacity of the LUN equal the user capacity of the expanded Group. Whether you can actually use the increased user capacity of the LUN depends on the operating system running on the servers connected to the storage system.

If the storage system supports metaLUNs, the **Expand LUN with RAID Group** feature is unavailable.

RAID group defragmentation

If you unbind and bind LUNs on a RAID Group, you may create gaps in the contiguous space across the Group's disks. This activity, fragmenting the RAID Group, leaves you with less space for new LUNs. You can defragment a RAID Group to compress these gaps and provide more contiguous free space across the disks. Defragmentation may also shorten file access time, since the disk read/write heads need to travel less distance to reach data.

Disk replacement

When a disk in a RAID Group is replaced or fails, the rebuild operation reconstructs the data on the replacement disk or hot spare one LUN at a time, starting with the first LUN.

RAID group properties

A RAID Group has the following properties:

RAID Group ID - Each RAID Group in a storage system has a unique ID, which is a hexadecimal number. The default ID is the smallest available number.

Expansion/defragmentation priority - Determines how fast expansion and defragmentation occurs. Values are Low, Medium, or High.

Automatically destroy - Enables or disables (default) the automatic dissolution of the RAID Group when the last LUN in that RAID Group is unbound.

Table 8-7 Default RAID group property values

Property	Value
RAID Group ID	Smallest available number
Expansion/Defragmentation Priority	Medium
Automatically Destroy	Cleared

Table 8-8 Maximum number of LUNs per RAID group

LUN RAID type	Maximum LUNs per RAID group
RAID 5, RAID 3*, RAID 1/0, RAID 1, RAID 0, and Disk	128
Hot Spare	Only 1 is allowed

* Multiple LUNs may be created from a RAID 3 Raid Group, but the highest bandwidth will be achieved with 1 to 4 threads of concurrent, large I/O.

RAID group characteristics

The following characteristics apply to RAID Groups and the LUNs on them:

- ◆ All LUNs in a RAID Group have the same RAID type
- ◆ Each LUN in a RAID Group can have a different element size where applicable
- ◆ Different SPs can own different LUNs in the same RAID Group

What next?

What you do next depends on whether the storage system supports RAID Groups.

For a RAID Group storage system – Continue to the next section, "Creating RAID groups".

For a non-RAID Group storage system – Go to the section, "Creating LUNs on a non-RAID group storage system" on page 8-21.

Creating RAID groups

Before you create LUNs in a RAID Group storage system, you must first create the RAID Group on which you will bind the LUN. You create a RAID Group using the **Create RAID Group** dialog box. You can create either:

- ◆ RAID Groups with disks and default property values that Manager selects
- ◆ RAID Groups with disks and property values that you select

Do not mix Fibre Channel and ATA (Advanced Technology Attachment) disks within the same RAID Group.

The table below lists the maximum number of RAID Groups allowed per storage system based on storage system type:

CX700, CX600	CX500	CX400, CX300	CX200	CX200LC
240	120	60	30	2

To create a RAID group

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the storage system on which you want to create a RAID Group, right-click, and then select **Create RAID Group**.
2. In the **Create RAID Group** dialog box, select a **RAID Group ID**.

You cannot change the ID of the RAID Group without destroying the RAID Group (and thus unbinding all its LUNs and losing their data) and then recreating it with the new ID.

For a description of each property in the dialog box, click **Help**.

3. To create a RAID Group with disks and default values that Manager selects, do the following:
 - a. Select automatic disk selection.
 - b. Select the number of disks that this RAID Group will contain.
 - c. Click **Apply** to create the RAID Group and close the dialog box.

Manager creates the RAID Group with disks it selects, and sets default values for the advanced properties.

If the storage system uses both ATA and Fibre Channel disks, Manager will initially try to configure the RAID Group using ATA disks. If the number of selected disks exceeds the number of available ATA disks, Manager will configure the RAID Group using Fibre Channel disks.

4. To create a RAID Group with disks and property values that you select, do the following:
 - a. Select manual disk selection.
 - b. Select the disks you want to include in this RAID Group.
 - c. Change any of the advanced property values for the RAID Group.
 - d. Click **Apply** to create the RAID Group and close the dialog box.

Manager creates the RAID Group with the disks you selected, and the values you assigned.

A newly created RAID Group with no bound LUNs displays a **Total Capacity** value of .000.

What next? When you have created the RAID Groups you want, continue to the next section to bind one or more LUNs on each of them.

Creating LUNs on RAID groups

When you bind a LUN on a RAID Group, you specify how much of the Group's user space (contiguous free space) you want the LUN to use. The LUN is distributed equally across all the disks in the RAID Group.

For example, a RAID 5 RAID Group with five 9-Gbyte disks provides 36 Gbytes of user space and 9 Gbytes of parity data. If you bind one 2-Gbyte LUN, you will have 34 Gbytes left for additional LUNs. You could bind 17 more 2-Gbyte LUNs using all the space in the RAID Group, or you could bind four more 2-Gbyte LUNs and four 5-Gbyte LUNs, leaving 6 Gbytes for future expansion. The maximum number of LUNs in a RAID Group is 128, and the maximum number in a Storage Group is 256. The maximum number of LUNs in a storage system depends on the storage system type, as shown in the following table:

CX700	CX600, CX500	CX400, CX300	CX200, CX200LC
2048	1024	512	256



CAUTION

Pre-FC4700 storage systems

Assign at least 2 Mbytes per RAID 3 LUN to the RAID 3 memory partition. If this partition does not have adequate memory for the LUN, you will not be able to bind it. Changing the size of the RAID 3 memory partition reboots the storage system. Rebooting restarts the SPs in the storage system, which terminates all outstanding I/O to the storage system.

CX-Series or FC4700-Series storage systems that support RAID 3 LUNs

Allocating the RAID 3 memory partition size is not required for RAID 3 RAID Groups and LUNs (the RAID 3 memory partition appears dimmed and is unavailable). If there will be a large amount of sequential read access to this RAID 3 LUN, it may be beneficial to enable read caching with prefetching for the LUN.

To create LUNs on a RAID group

You create a LUN using the **Bind LUN** dialog box.

If no LUNs exist on a storage system connected to a NetWare server, refer to the Release Notice for the NetWare Navisphere Agent for information on how to bind the first LUN.

If you are binding LUNs on a storage system connected to a Solaris server, and no failover software is installed, refer to the *Storage System Host Utilities for Solaris Administrator's Guide* (069001140) for information on how to bind the first LUN.

If the LUNs you are binding reside on a storage system connected to a VMware ESX Server, and these LUNs will be used with layered applications such as SnapView, configure the LUNs as raw device mapping volumes and set the state to physical compatibility mode.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the storage system on which you want to bind a LUN, right-click, and then select **Bind LUN**.
2. In the **Bind LUN** dialog box, select a RAID type for the new LUN.
3. In the **RAID Group for new LUN** list, click the ID of the RAID Group on which you want to bind the new LUNs.

The list displays only those RAID Groups available for the selected RAID type.

If there are no RAID Groups for the RAID type specified, click **New** to create one.

4. Set any of the following LUN properties:

For a description of each property, click **Help** in the dialog box or refer to "About LUNs" on page 8-2.

- **LUN ID** (default is smallest available number)

Once the LUN is bound, you cannot change the LUN ID without unbinding the LUN (and thus losing its data) and then rebinding it with the new ID.

If this CLARiiON storage system is connected to a Celerra workstation, you must manually assign the LUN ID for any new LUNs. The LUN ID must be 16 or greater.

- Rebuild priority

- Verify priority
- Enable read cache, and write caching

A LUN that is bound with read caching enabled uses caching only if the read cache is enabled for the SP that owns it (page 7-15). Similarly, a LUN bound with write caching enabled uses caching only if storage-system write caching is enabled (page 7-15).

- Number of LUNs to bind with the specified values
- Default owner for the new LUNs

The default owner is the SP that assumes ownership of the LUN when the storage system is powered up. If the storage system has two SPs, you can choose to bind some LUNs using one SP as the default owner and the rest using the other SP as the default owner, or you can select **Auto**, which tries to divide the LUNs equally between SPs. The primary route to a LUN is the route through the SP that is its default owner, and the secondary route is through the other SP.

If you do not specifically select one of the default owner values, Manager assigns default LUN owners according to RAID Group IDs as follows:

RAID group IDs	Default LUN owner
Odd-numbered	SP A
Even-numbered	SP B

- LUN size (MB, GB or Block Count)
 - Enable auto assign
 - Alignment offset
 - No initial verify
5. Click **Apply** to bind the LUN and close the dialog box.

Binding LUNs may take as long as two hours. Some storage systems have disks that have been preprocessed at the factory to speed up binding. You can determine the progress of a bind operation from **Percent Bound** on the **LUN Properties** dialog box (page 13-16).

The **RAID Group - Partitions** tab will display incorrect data while LUNs are being bound on the RAID Group. The data will be correct once the LUNs are completely bound.

6. If the LUN you just bound is on a CX-Series or FC-Series storage system, follow the procedure below to eliminate any latent software errors before using the LUN.
 - a. When the bind operation is complete, use the following Navisphere CLI command to determine which SP owns the LUN.

navicli -h *hostname* getlun *lun* -owner

where

hostname specifies the IP address or network name of an SP in the storage system.

lun Specifies the logical unit number of the LUN.

- b. Start a background verify and set the sniffer rate for the LUN with the following Navisphere CLI Command:

This step is not required for systems running storage system software version 02.05 or later because background verifies are done automatically.

navicli -h *hostname* setsniffer *lun* 1 -bv -bvtime ASAP -snrate *rate*

where

hostname specifies the IP address or network name of an SP in the storage system.

lun Specifies the logical unit number of the LUN.

rate Specifies the sniff rate. The rate you should use varies with the capacity of the disks in the LUN, as shown in the table below. The default rate is 30.

Disk size	Recommended	
	Sniffer rate	Maximum sniffer rate
36 Gbytes	20	30
73 Gbytes	10	20
148 Gbytes	5	10
250 Gbytes	5	5
320 Gbytes	5	5

If you want to use a higher rate than the maximum, consult technical support before setting it to the desired rate.

What next?

What you do after you have created all the LUNs you want depends on whether the storage system is unshared or shared.

Unshared storage systems - Reboot each server connected to the storage system to make the LUNs in the storage system visible to the server.

Shared storage system - Go to Chapter 9, "Setting Up Access Logix and Creating Storage Groups", to create Storage Groups containing the LUNs you bound.

Binding a new LUN with the exact size of an existing LUN

If you want to bind a new LUN with the exact size of an existing LUN, do the following:

1. In the **Storage** tree, navigate to the LUN whose size you want you duplicate, right-click, and then select **Properties**.
2. Record the value shown in **User Blocks** and then when you bind the new LUN, enter this value in **LUN Size**, and select **Block Size**.

Creating LUNs on a non-RAID group storage system

You can create either:

- ◆ LUNs with disks that Manager selects and default property values (page 8-10), or
- ◆ LUNs with disks that you select and property values that you set



CAUTION

Before you bind a RAID 3 LUN, you must assign memory for it to the RAID 3 memory partition. If this partition does not have adequate memory for the LUN, you will not be able to bind it. Changing the size of the RAID 3 memory partition reboots the storage system.

To bind LUNs on a non-RAID group storage system

If no LUNs exist on a storage system connected to a NetWare server, refer to the Release Notice for the NetWare Navisphere Agent for information on how to bind the first LUN.

If you are binding LUNs on a storage system connected to a Solaris server, and no failover software is installed, refer to the Release notes for the Solaris Navisphere Agent for information on how to bind the first LUN.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the storage system on which you want to bind a LUN, right-click, and then select **Bind LUN**.
2. In the Bind LUN dialog box, set the following values:
 - LUN ID (default is smallest available number).

Once the LUN is bound, you cannot change the LUN ID without unbinding the LUN (and thus losing its data) and then rebinding it with the new ID.

- RAID type (this sets the RAID type for the new LUNs RAID Group).

3. To bind a LUN with disks that Manager selects, and default property values, do the following:
 - a. Complete step 2, and then, select the number of disks that you want to include in this LUN.

- b. Click **Apply** to bind the LUN and close the dialog box.

Manager binds a LUN or LUNs with disks it selects and assigns default values for all advanced properties.

4. To bind a LUN with disks that you select, and advanced property values, do the following:
 - a. Complete step 2, and then, under **Disk Selection**, select either **Manual** or **Automatic**.

Automatic lets you create multiple LUNs with the same RAID type and same number of disks in one operation. **Manual** lets you create only one LUN with the specified properties.

- b. Set or change any of the following values:

- **Default Owner (SP A, SP B or Auto)**

If the storage system has only one SP, **Auto** assigns SP A as the owner of all new LUNs. If the storage system has two SPs, it distributes the new LUNs between the two SPs. In so doing, it takes into account any existing LUNs owned by the SPs. As a result, either both SPs end up with the same number of LUNs, or one SP ends up with one more LUN than the other.

- **Number of Disks** to bind into this LUN with the specified values.
- **Rebuild Priority**
- **Verify Priority**
- **Enable Read Cache**, and **Enable Write Cache**

A LUN that is bound with read caching enabled uses caching only if the read cache is enabled for the SP that owns it (page 7-15). Similarly, a LUN bound with write caching enabled uses caching only if storage-system write caching is enabled (page 7-15).

- **Enable Auto Assign**

- c. Click **Apply** to bind the LUN, and then click **Cancel** to close the dialog box.

Manager binds a LUN or LUNs with the disks you selected and the values you assigned to any Advanced LUN properties.

Binding LUNs may take as long as two hours. Some storage systems have disks that have been preprocessed at the factory to speed up binding. You can determine the progress of a bind operation from **Percent Bound** on the **LUN Properties** dialog box (page 13-16).

What next?

- ◆ If you want to expand the user capacity of any LUNs before assigning them to Storage Groups, go to Chapter 10, "Expanding LUN Capacity With MetaLUNs".
- ◆ If not, then go to Chapter 9, "Setting Up Access Logix and Creating Storage Groups".

The Access Logix™ feature lets you enable data access and create Storage Groups on shared storage systems. A Storage Group is a collection of one or more LUNs or metaLUNs to which you connect one or more servers. A server can access only the LUNs in the Storage Group to which it is connected. It cannot access LUNs assigned to other servers. In other words, the server sees the Storage Group to which it is connected as the entire storage system.

For security reasons, we recommend that you place all storage systems in a domain before you configure them.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes the following:

- ◆ Setting the storage-system data access property 9-2
- ◆ Adding LUNs to a storage group 9-4
- ◆ Connecting a server to a storage group 9-5
- ◆ Verifying server connections to a storage group 9-6
- ◆ Verifying or editing device information in the host agent configuration file (pre-FC4700 storage systems) 9-9

Setting the storage-system data access property

The storage-system data access property - **access control enabled** - is available for shared storage systems only (those that support Storage Groups). This section describes data access control and Storage Groups and how to enable data access control for a storage system.

Data access control and storage groups

A shared storage system provides a data access control feature that lets you restrict the servers that can read and write to specific LUNs on the storage system. This feature is implemented using Storage Groups.

A Storage Group is a collection of one or more LUNs to which you connect one or more servers. A server can access only those LUNs in the Storage Groups to which it is connected. In other words, a server sees the Storage Groups to which it is connected as the entire storage system.

If you do *not* enable data access control for a storage system, the following occurs:

- ◆ You cannot create Storage Groups.
- ◆ Each server connected to the storage system has access to all LUNs on the storage system.

Enabling data access control for a storage system

If you want to create Storage Groups on a storage system, you must enable data access control for the storage system.

Important Access Logix must be installed on the storage system in order to enable data access control. If the **Data Access** section is not visible in the **Storage Access** tab, Access Logix is not installed on the storage system.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage-system icon for which you want to enable data access control, right-click, and then select **Properties**.
2. In the **Storage Access** tab, select the **Access Control Enabled** check box to enable access control.

Once you enable data access control for a storage system, Manager does not let you disable it. Contact Customer Service if you need to disable data access

Creating storage groups

For security reasons, we recommend that you place all storage systems in a domain before you configure the storage system.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

Before you can create Storage Groups on a storage system, you must have enabled data access control for the storage system, as described in "Setting the storage-system data access property" on page 9-2.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system on which you want to create the **Storage Group**, right-click, and then select **Create Storage Group**.
2. Assign a name to your Storage Group, and click **OK** or **Apply** to save your changes.

The following table lists the maximum number of Storage Groups allowed per storage system based on storage system type:

CX700	CX600, CX500	CX400, CX300	CX200	CX200LC
512	256	128	30	2

Adding LUNs to a storage group

Before you can perform I/O to a Storage Group, you must add LUNs to a Storage Group. The maximum number of LUNs in a Storage Group is 256.

If you will be adding a LUN to more than one Storage Group, all servers connected to those Storage Groups must be part of a cluster configuration.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group to which you want to add LUNs, right-click, and then choose **Select LUNs**.
2. In the **Storage Group Properties - LUNs** tab, select the LUN or LUNs that you want to add to this Storage Group, and move them into **Selected LUNs**.

In **Selected LUNs**, you can manually assign a **Host ID** to any newly added LUNs, or you can let the system assign the ID. Once you have assigned the Host ID to the LUN, the only way to change it is to remove the LUN from the Storage Group, and then reassign to a Storage Group as a new LUN.

- a. To manually assign a Host ID, click in the Host ID column, and select an ID from the list of available IDs.
 - b. To let the system assign the Host ID, leave the Host ID column empty.
3. Click **OK** or **Apply** to save your changes.

Connecting a server to a storage group

You must connect a server to a Storage Group so that the server can perform I/O to the LUNs in the Storage Group. This server can access only the LUNs in the Storage Group to which it is connected.

A server can connect to only one Storage Group per storage system at a time. If you try to connect a server, that already connects to a Storage Group, to another Storage Group in the same storage system, the server will disconnect from the first Storage Group. When you disconnect a server from a Storage Group, and then connect it to a different Storage Group, the server can no longer perform I/O to the first Storage Group.

If the storage system is connected to a VMware ESX Server, use Navisphere Manager to connect the ESX Server to the LUNs in the Storage Group; rescan the bus at the ESX Server level; power down any virtual machines to which you will assign the LUN; use the ESX Server console to assign the LUNs to the virtual machines; power up the virtual machine.

When you connect a server to a Storage Group, the server is

- ◆ Connected to the Storage Group through each server HBA port (initiator) that is connected to the storage system.

In other words, if the server has two HBA ports and each is connected to one storage-system SP, the server has two connection paths to the Storage Group.

- ◆ Disconnected from any other Storage Group in the storage system to which it is connected.

Navisphere lets you connect multiple servers to the same Storage Group, but when you do, Navisphere displays a confirmation message requiring that you click **Yes** to continue or **No** to cancel the operation. You should do this only if the following conditions are met:

- ◆ The servers are running the same operating system.
- ◆ The operating system either
 - supports multiple sharing of the same LUN, or
 - has layered software (such as Microsoft Cluster) that supports multiple servers sharing the same LUN.

In a cluster environment, you must connect each server in the cluster to the same Storage Group.

In order to display accurate server/LUN information about a LUN that is shared between servers in a cluster, right-click the shared LUN, and click **Update Host Information**. Navisphere displays updated information about the server that currently controls the shared LUN in the **LUN Properties - Hosts** tab.

Connecting a server to a storage group

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group to which you want to connect a server, right-click, and then select **Connect Hosts**.
2. In **Available Hosts**, select the server or servers that you want to connect to this Storage Group, and move them into the **Hosts to be Connected** list.
3. Click **OK** or **Apply** to save your changes.

What next?

Continue to the next section to verify the connections to the Storage Groups you just created.

Verifying server connections to a storage group

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system for which you want to verify Storage Group connections, right-click, and then select **Connectivity Status**.

The **Connectivity Status** dialog box includes an entry for each HBA connected to the selected storage system that is currently logged in or permanently registered.

An HBA can have a logged-in state of active or inactive. When an HBA's logged-in state is active:

- The HBA is powered up
- There is a connection from the server to the storage system
- The FLARE software recognizes the HBA.

When an HBA's logged state is inactive, make sure that:

- The HBA port is working.
- The switch is working, and it is configured to connect the HBA port to the SP. Refer to the switch documentation for information on the switch and how to configure it.
- The cables between the HBA port and the switch, and between the switch and the SP, are fastened securely.

If you determine that there is a valid and operational path from the HBA to the SP, then either the SP or the HBA to which it was connected was replaced, and the information about the connection between the SP and the HBA was never removed from the storage system's persistent memory.

Whenever a storage-system server is rebooted, the Agent scans the network for HBA port connections to storage systems. When it finds a connection, it sends information about the connection to the SP.

The SP stores this information in the storage system's persistent memory on the database disks. This information remains in this memory until you issue a CLI **port** command to remove it. Refer to the Agent and CLI manual for information on the **port** command.

An HBA is permanently registered when the Navisphere Agent recognizes the HBA. When all HBAs belonging to a server are registered, you can connect this server to Storage Groups.

2. When you have verified the server connections to each Storage Group in the storage system, reboot each server connected to the storage system.

Rebooting makes the LUNs in the Storage Group connected to the server visible to the server.

What next? **For pre-FC4700 storage systems** - Continue to the next section "Verifying or editing device information in the host agent configuration file (pre-FC4700 storage systems)" on page 9-9.

For CX-Series or FC4700-Series storage systems

- ◆ You must make the LUNs available to the operating system as described in the documentation for your specific operating system, or refer to the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems* available on Powerlink.
- ◆ If you will be using data replication software such as SnapView, incremental SAN Copy, or MirrorView / A, you must configure the reserved LUN pool (refer to Chapter 11, "Managing the Reserved LUN Pool for LUN Replication"). If not, then proceed to Chapter 12, "Setting Up and Using Event Monitor".

Verifying or editing device information in the host agent configuration file (pre-FC4700 storage systems)

Whenever you create one or more LUNs on a storage system, you should verify or edit the device information in the Host Agent configuration file on the server. This section describes how to do this for the following servers:

- ◆ AIX server (this page)
- ◆ HP-UX server (this page)
- ◆ Linux server (this page)
- ◆ NetWare server (this page)
- ◆ Solaris server (page 9-10)

Verifying or editing host agent device information on an AIX, HP-UX, Linux, or NetWare or Windows server

1. Open the Host Agent configuration file.
 - If the file contains a **device** entry for an existing LUN on each SP in the storage system, do not change anything.
 - If the file does *not* contain a **device** entry for an existing LUN on each SP in the storage system, add one for each SP.
2. If you did not change anything, just close the file; if you changed the file, save it.
3. Stop and then start the Agent.

What next?

AIX, HP-UX, Linux, NetWare or Windows on the server views the LUNs in a storage system as identical to standard single disk drives. For AIX, HP-UX, Linux, NetWare or Windows to use the LUNs, you must make them available to the operating system as described in the operating system documentation. Refer to documentation for your specific operating system, or refer to the *EMC Installation Roadmap for CX-Series, AX-Series, and FC-Series Storage Systems* available on Powerlink.

Verifying or editing agent device information on a Solaris server

How you edit the Host Agent configuration file depends on whether the storage system had any bound LUNs before you created LUNs. Follow the appropriate procedure below for your situation.

No failover software and no bound LUNs before you created LUNs

If you are binding LUNs in a storage system connected to a Solaris server and *no* LUNs exist on the storage system, edit the device information in the Host Agent configuration file on each server connected to the storage system in one of the following ways:

1. Open the Agent configuration file and enter the following line entry:

device auto auto "auto"
2. Save the Agent configuration file before stopping and then starting the Agent.

or
1. Open the Agent configuration file and add a **clspn** entry for each SP in the storage system.

For information on **clspn** entries, refer to the Agent manual for Solaris environments.
2. Comment out (insert a # before) any **device** entries for the SPs in the storage system.
3. Save the Agent configuration file before stopping and then starting the Agent.

What next?

The Solaris operating system running on the server views the LUNs in a storage system as identical to standard single disk drives. For Solaris to use the LUNs, you must make them available to the operating system. Refer to documentation for your specific operating system, or refer to the *EMC Installation Roadmap for CX-Series, AX-Series, and FC-Series Storage Systems* available on Powerlink.

Bound LUNs before you created LUNs

1. Open the Agent configuration file and comment out (insert a # sign before) any **clsp device** entries for SPs in the storage system.
 - If the file contains a **device** entry for an existing LUN on each SP in the storage system, do not change anything.
 - If the file does *not* contain a **device** entry for an existing LUN on each SP in the storage system, add one for each SP.
2. If you did not change anything, just close the file; if you changed the file, save the file.
3. Stop and then start the Agent.

What next?

Solaris on the server views the LUNs in a storage system as identical to standard single disk drives. For Solaris to use the LUNs, you must make them available to the operating system as described in the Solaris documentation.

This chapter explains how to increase LUN user capacity with metaLUNs.

Major sections in the chapter are

- ◆ MetaLUN terminology 10-2
- ◆ MetaLUNs overview 10-3
- ◆ Expanding a LUN 10-5
- ◆ Expanding a metaLUN 10-9
- ◆ Using the expand storage wizard to expand LUN capacity 10-12
- ◆ Displaying and changing MetaLUN properties 10-17
- ◆ Monitoring expansion status 10-19
- ◆ Destroying a metaLUN 10-19
- ◆ Using the added capacity of an expanded SnapView,
MirrorView / A, MirrorView / S or SAN Copy LUN 10-20

MetaLUN terminology

Term	Meaning
base LUN	A single LUN that is expanded into a metaLUN.
expansion LUNs	LUNs that are used to expand the capacity of the base LUN or a metaLUN.
component	Grouping of one or more LUNs that is part of a metaLUN. If more than one LUN is included in the grouping, the LUNs will be striped together.
metaLUN	A storage component consisting of 2 or more LUNs whose capacity can grow dynamically by adding additional LUNs.
stripe expansion	When expanding a single FLARE LUN, creates a metaLUN with one component which includes the base LUN and any LUNs that you add. Any data on the base LUN is re-striped across the base LUN and the new LUNs. When expanding a metaLUN, re-stripes any data on the last component of a metaLUN across this component and any new LUNs that you add to the component. The new LUNs become part of the metaLUN component.
concatenate expansion	Creates a new metaLUN component that includes the new expansion LUNs and appends this component to the existing LUN or metaLUN as a single, separate striped component. There is no striping between the existing storage and the new component.

MetaLUNs overview

A metaLUN is a type of LUN whose maximum capacity can be the combined capacities of all the LUNs that compose it. The metaLUN feature lets you dynamically expand the capacity of a single LUN (base LUN) into a larger unit called a metaLUN. You do this by adding LUNs to the base LUN. You can also add LUNs to a metaLUN to further increase its capacity. Like a LUN, a metaLUN can belong to a Storage Group, and can participate in SnapView, MirrorView and SAN copy sessions. MetaLUNs are supported only on CX-Series storage systems.

A metaLUN may include multiple sets of LUNs and each set of LUNs is called a component. The LUNs within a component are striped together and are independent of other LUNs in the metaLUN. Any data that gets written to a metaLUN component is striped across all the LUNs in the component. The first component of any metaLUN always includes the base LUN. The number of components within a metaLUN and the number of LUNs within a component depend on the storage system type. The following table shows this relationship:

Storage system type	LUNs per metaLUN component	Components per metaLUN
CX700, CX600	32	16
CX500, CX500i, CX400	32	8
CX300, CX200	16	8

You can expand a LUN or metaLUN in two ways — stripe expansion or concatenate expansion. A stripe expansion takes the existing data on the LUN or metaLUN you are expanding, and restripes (redistributes) it across the existing LUNs and the new LUNs you are adding. The stripe expansion may take a long time to complete. A concatenate expansion creates a new metaLUN component that includes the new expansion LUNs, and appends this component to the existing LUN or metaLUN as a single, separate, striped component. There is no restriping of data between the original storage and the new LUNs. The concatenate operation completes immediately.

During the expansion process, the server is able to process I/O to the LUN or metaLUN, and access any existing data. It does not, however, have access to any added capacity until the expansion is complete. When you can actually use the increased user capacity of the

metaLUN depends on the operating system running on the servers connected to the storage system.

Figure 10-1 depicts a LUN using stripe expansion to add two LUNs. The result is a metaLUN with a single component consisting of all three LUNs with the data on the base LUN re-striped across all three LUNs.




Figure 10-1 Sample stripe expansion

Figure 10-2 depicts a LUN using concatenate expansion to add two LUNs. The result is a metaLUN with two components. The first component is the base LUN. The second component consists of two new LUNs striped together, but not striped to the first component.




Figure 10-2 Sample concatenate expansion

EMC recommends that you do not concatenate LUNs of different RAID types. Do this only in an emergency situation when you need to add capacity to a LUN and you do not have LUNs of the same RAID type or the disk capacity to bind new ones. Concatenating metaLUN components with a variety of RAID types could impact the performance of the resulting metaLUN. Once you expand a LUN, you cannot change the RAID type of any of its components without destroying the metaLUN. Destroying a metaLUN unbinds all LUNs in the metaLUN, and therefore causes data to be lost.

What next Depending on the user options that you specify when you create or expand a metaLUN, its resulting geometry and performance can vary greatly.

- ◆ If you are expanding a LUN, see "Expanding a LUN" on page 10-5.
- ◆ If you are expanding a metaLUN, see "Expanding a metaLUN" on page 10-9.

Expanding a LUN

This section describes the restrictions that apply as well as what happens when you expand a LUN (the base LUN).

Restrictions

The following restrictions apply when you expand a LUN:

- ◆ You cannot expand a private or transitioning LUN.
- ◆ Each LUN within a metaLUN must be either a redundant (RAID 5, RAID 3, RAID 1, RAID 1/0) or non-redundant (RAID 0, Disk) RAID type.
- ◆ All disks in the LUNs must be the same type — ATA or Fibre Channel.
- ◆ Expansion LUNs must not be
 - Part of a Storage Group.
 - Private or transitioning.
- ◆ If the LUN you wish to expand is participating in a SnapView (snapshot or clone), MirrorView, or SAN Copy session, you will be able to add additional capacity to the LUN, but the wizard will *not* let you modify the user capacity at this time. The wizard will automatically select **Current Capacity** for you. This keeps the available user capacity of the LUN the same as it was before adding LUNs.

In addition, if the LUN meets one of the following conditions, you must complete the associated task before you can expand it:

- If the LUN is participating in a MirrorView session and it is in the synchronizing state, you must wait for the LUN to finish synchronizing before you can expand it.

- If the LUN is participating in a SAN Copy session, and it is actively being copied, you must stop the session or wait for the session to complete.

In order to change the user capacity of the expanded LUNs, you must wait for the expansion to complete and then follow the instructions in "Using the added capacity of an expanded SnapView, MirrorView / A, MirrorView / S or SAN Copy LUN" on page 10-20.

- ◆ If the LUN you are expanding is connected to a VMware ESX Server, it must be configured as a raw device mapping volume, set to compatibility mode. In order for the ESX Server and virtual machines to recognize the added capacity, you must rescan the bus at the ESX Server level, and optionally, at the virtual machine level.

If the LUN you want to expand is not configured as a raw device mapping volume, use the VMware **vmkfstools** utility to reconfigure it. For information on using this utility, refer to *Managing Raw Device Mappings* technical note found at http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf.

Results

When you expand the capacity of the base LUN, the following happens:

- ◆ The new metaLUN assumes the LUN ID, World Wide Name (WWN) and LUN name of the base LUN.
- ◆ The software assigns a new LUN ID and WWN to the base LUN and makes it a private LUN.
- ◆ The software assigns new LUN IDs to any expansion LUNs, makes them private LUNs, and, if necessary, trespasses them to the SP that owns the base LUN. The expansion LUNs retain their original WWNs and LUN names.
- ◆ In the **Storage** tree, Manager places an icon for each newly created metaLUN below the SP to which it belongs.
- ◆ You will have access to any existing data on the base LUN during the expansion process.
- ◆ You will not have access to the metaLUN's additional capacity until the expansion process is complete. A concatenate expansion always completes immediately, but a stripe expansion will take longer if the base LUN has any existing data.

Table 10-1 describes what happens when you expand a LUN with and without existing data that is either in or out of a Storage Group:

If the LUN you want to expand is not in a Storage Group, and has existing data, and you are using stripe expansion, Manager gives you the option of preserving the data on the LUN or overwriting it. This option applies only to stripe expansion for LUNs that are not in a Storage Group.

Table 10-1 The effects of expanding a LUN

Base LUN description	Expansion type	Expansion description	Expansion rate
In a Storage Group with existing data	Stripe	The existing data is restriped across the base LUN and any expansion LUNs. All LUNs become part of a single metaLUN component. All expansion LUNs must be the same size and RAID type as the base LUN.	Slow
	Concatenate	The existing data is preserved on the base LUN. The base LUN becomes the first component in the metaLUN and the new storage is appended to the base LUN component as a single, striped component. No striping occurs between the base LUN and the expansion LUNs. LUNs within the second metaLUN component must all be the same size and RAID type, but this component can differ in size and RAID type from LUNs in other components ^a .	Immediate
In a Storage Group with <i>no</i> existing data	Stripe	The new storage is combined with the base LUN into a single, striped component. All expansion LUNs must be the same size and RAID type as the base LUN.	Immediate
	Concatenate	The new storage is appended to the base LUN component as a single, striped unit (metaLUN component). No striping occurs between the base LUN component and the expansion LUNs. LUNs within the metaLUN component must all be the same size and RAID type, but they do not need to be the same as LUNs in any other metaLUN component ^a .	Immediate
<i>Not</i> in a Storage Group with existing data	Stripe (preserve existing data)	The existing data is restriped across the base LUN and any expansion LUNs. All LUNs become part of a metaLUN component. All expansion LUNs must be the same size and RAID type as the base LUN.	Slow
	Stripe (do not preserve existing data)	Data on the base LUN is discarded. The new storage combines with the base LUN into a single, striped component. All expansion LUNs must be the same size and RAID type as the base LUN.	Immediate

Table 10-1 The effects of expanding a LUN (*continued*)

Base LUN description	Expansion type	Expansion description	Expansion rate
<i>Not</i> in a Storage Group with existing data (cont)	Concatenate	The existing data is preserved on the base LUN. The new storage is appended to the base LUN component as a single, striped unit (metaLUN component). No striping occurs between the base LUN component and the LUNs in the new component. LUNs within the new metaLUN component must all be the same size and RAID type, but can be different from LUNs in any other metaLUN component ^a .	Immediate
<i>Not</i> in a Storage Group with <i>no</i> existing data	Stripe	The new storage is combined with the base LUN into a single, striped component. All expansion LUNs must be the same size and RAID type as the base LUN.	Immediate
	Concatenate	The new storage is appended to the base LUN component as a single, striped unit (metaLUN component). No striping occurs between the base LUN component and the LUNs in the new component. LUNs within the new metaLUN component must all be the same size and RAID type, but can be different from LUNs in any other metaLUN component ^a .	Immediate

- a. EMC recommends that you do not expand LUN capacity by concatenating LUNs of different RAID types. We recommend that you do this only in an emergency situation when you need to add capacity to a LUN and you do not have LUNs of the same RAID type or the disk capacity to bind new ones. Concatenating metaLUN components with a variety of RAID types could impact the performance of the resulting metaLUN. Once you expand a LUN, you cannot change the RAID type of any of its components without destroying the metaLUN. Destroying a metaLUN unbinds all LUNs in the metaLUN, therefore losing all data.

Expanding a metaLUN

This section describes the restrictions that apply as well as what happens when you expand a metaLUN.

Restrictions

The following restrictions apply when you expand a metaLUN:

- ◆ You can stripe expand only the *last* component of a metaLUN.
- ◆ You cannot expand a private or transitioning metaLUN.
- ◆ Each LUN within the metaLUN must be either a redundant (RAID 5, RAID 3, RAID 1, RAID 1/0) or non-redundant (RAID 0, Disk) RAID type.
- ◆ All disks in the LUNs that compose the metaLUN must be the same type - ATA or Fibre Channel.
- ◆ Expansion LUNs must not be
 - Part of a snapshot, clone, mirror or SAN Copy session.
 - Part of a Storage Group
 - Private or transitioning.
- ◆ If the metaLUN you wish to expand is participating in a SnapView (snapshot or clone), MirrorView, or SAN Copy session, you will be able to add additional capacity to the metaLUN, but the wizard will *not* let you modify the user capacity at this time. The wizard will automatically select **Current Capacity** for you. This keeps the available user capacity of the metaLUN the same as it was before adding LUNs.

In addition, if the metaLUN meets one of the following conditions, you must complete the associated task before you can expand it:

- If the metaLUN is participating in a MirrorView session and it is in the synchronizing state, you must wait for the metaLUN to finish synchronizing before you can expand it.
- If the metaLUN is participating in a SAN Copy session, and it is actively being copied, you must stop the session or wait for the session to complete.

In order to change the user capacity of the expanded metaLUNs, you must wait for the expansion to complete and then follow the instructions in "Using the added capacity of an expanded SnapView, MirrorView / A, MirrorView / S or SAN Copy LUN" on page 10-20.

- ◆ If the metaLUN you are expanding is connected to a VMware ESX Server, it must be configured as a raw device mapping volume, set to compatibility mode. In order for the ESX Server and virtual machines to recognize the added capacity, you must rescan the bus at the ESX Server level, and optionally, at the virtual machine level.

If the LUNs that comprise the metaLUN you want to expand are not configured as a raw device mapping volumes, use the VMware **vmkfstools** utility to reconfigure them. For information on using this utility, refer to the Managing *Raw Device Mappings* technical note found at http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf.

Results

When you expand the capacity of a metaLUN, the following happens:

- ◆ The software assigns new LUN IDs to any expansion LUNs, makes them private LUNs, and, if necessary, trespasses them to the SP that owns the base LUN. The expansion LUNs retain their original WWNs and LUN names.
- ◆ In the **Storage** tree, under the icon for the metaLUN, Manager adds icons for each LUN and new metaLUN component that you add to the metaLUN.
- ◆ You will have access to any existing data on the metaLUN during the expansion process.
- ◆ You will not have access to the metaLUN's additional capacity until the expansion process is complete. A concatenate expansion always completes immediately, but a stripe expansion will take longer if the base LUN has any existing data.

Table 10-2 describes what happens when you expand various metaLUN configurations

The software always assumes that metaLUNs contain valid data. Therefore, data is always preserved when you expand a metaLUN, regardless of whether or not the metaLUN is in a Storage Group.

Table 10-2 The effects of expanding a metaLUN

Beginning metaLUN configuration	Expansion type	Expansion description	Expansion rate
MetaLUN with only <i>one</i> metaLUN component containing the base LUN and one or more additional LUNs.	Stripe	The expansion LUNs are added to the metaLUN component. Any existing data on the metaLUN component is restriped across the existing LUNs and any expansion LUNs. All expansion LUNs must be the same size and RAID type as the base LUN.	Slow
	Concatenate	The existing data is preserved on the original metaLUN component. Expansion LUNs are added as a single, striped unit (metaLUN component). No striping occurs between the original component and the new component. LUNs within the new metaLUN component must all be the same size and RAID type, but can be different from the original component ^a .	Immediate
MetaLUN with more than one metaLUN component. (You can expand only the last component of a metaLUN)	Stripe	Any data on the last component of the metaLUN is restriped across the existing LUNs in the component and any expansion LUNs. All expansion LUNs must be the same size and RAID type as the LUNs in the component you are expanding.	Slow
	Concatenate	The existing data is preserved on the original metaLUN. Expansion LUNs are added as a single, striped unit (metaLUN component) to the metaLUN. No striping occurs between the original metaLUN components and the new component. LUNs within the new metaLUN component must all be the same size and RAID type, but can be different from the LUNs in other components ^a .	Immediate

- a. EMC recommends that you do not expand LUN capacity by concatenating LUNs of different RAID types. We recommend that you do this only in an emergency situation when you need to add capacity to a LUN and you do not have LUNs of the same RAID type or the disk capacity to bind new ones. Concatenating metaLUN components with a variety of RAID types could impact the performance of the resulting metaLUN. Once you expand a LUN, you cannot change the RAID type of any of its components without destroying the metaLUN. Destroying a metaLUN unbinds all LUNs in the metaLUN, therefore losing all data.

Using the expand storage wizard to expand LUN capacity

The **Expand Storage Wizard** is the Manager tool that lets you dynamically expand the capacity of a LUN or metaLUN.

The software allows only 4 expansions per storage system to be running at the same time. Any additional requests for expansion are added to a queue, and when one expansion completes, the first one in the queue begins.

Starting the Expand Storage Wizard

The following steps start the **Expand Storage** wizard:

1. In the **Enterprise Storage** window, navigate to the LUN, metaLUN or metaLUN component that you want to expand (this object must reside on a storage system that supports metaLUNs).
2. Right-click the LUN or metaLUN and click **Expand**, or right-click the *last* component in the metaLUN and click **Add LUNs**.

You cannot expand a private LUN or metaLUN. If you select a private LUN, and click **Expand**, Manager displays an error message.

You can expand only the last component of a metaLUN. If you select a component other than the last one in the metaLUN, and click **Add LUNs**, Manager displays an error message.

The **Expand Storage** wizard opens.

The **Start Expand Storage Wizard** screen provides a brief overview of the steps for expanding a LUN, metaLUN or a metaLUN component.

3. Read the wizard summary and click **Next** to continue.

Selecting storage expansion type

The **Select Expansion Type** screen lets you select stripe or concatenate expansion. The graphic on the left side of the screen shows the difference between the two types of expansion. This screen displays only if you select a LUN or metaLUN to expand. It does not display if you select a metaLUN component since metaLUN components always use stripe expansion. Stripe is the default expansion type.

Select the expansion type and click **Next**.

Stripe expansion is the default.

- ◆ See "Expanding a LUN" on page 10-5.
- ◆ See "Expanding a metaLUN" on page 10-9.

Selecting expansion LUNs

Since you *cannot* use private LUNs or a LUNs that are in a Storage Group, the **Select Unused LUNs** wizard screen does not list private LUNs or LUNs in Storage Groups. You may need to bind new LUNs (see Chapter 8, "Creating LUNs and RAID Groups").

The **Select Unused LUNs** screen lets you select the LUNs you want to use to expand the capacity of a LUN, metaLUN or metaLUN component. The **Available LUNs** list displays only those LUNs that are appropriate and available for the selected expansion type.

For stripe expansions

- ◆ All listed LUNs are the same size and RAID type as the LUN, metaLUN or metaLUN component you are expanding.
- ◆ All listed LUNs are the same disk type (ATA or Fibre Channel) as the LUN, metaLUN or metaLUN component you are expanding.

For concatenate expansions

- ◆ All listed LUNs are the same redundancy type as the LUN, metaLUN or metaLUN component you are expanding, but they may not all be the same size and RAID Type.

Within a metaLUN component, all LUNs must be the same size and RAID Type, but can be different from LUNs in other components.

- ◆ All listed LUNs are the same disk type (ATA or Fibre Channel) as the LUN, metaLUN or metaLUN component you are expanding.

To select expansion LUNs, do the following:

1. In the **Available LUNs** list, select the LUNs you want to add, and click the right arrow.

As you select LUNs, the value for **Maximum user capacity** changes.

2. When you have selected all the LUNs you want to add, click **Next**.

Specifying metaLUN capacity

The **Specify MetaLUN Capacity** wizard screen lets you specify the user capacity for the metaLUN.

If the LUN or metaLUN you are expanding is participating in a SnapView (snapshot or clone), MirrorView, or SAN Copy session, you will not be able to change the user capacity at this time. The wizard automatically selects **Current Capacity** for you. This keeps the available user capacity of the metaLUN the same as it was before adding LUNs. In order to change the user capacity, you must wait for the expansion to complete and then follow the instructions in the Administrators Guide for SnapView, MirrorView or SAN Copy.

In **User Capacity**, select the unit of measure and capacity for the metaLUN, where

- ◆ **MB, GB, TB and Blocks** is megabytes, gigabytes, terabytes, or blocks, respectively.
- ◆ **Current capacity** keeps the available user capacity of the metaLUN the same as it was before adding LUNs. You cannot use any added capacity until you change the capacity in the **MetaLUN Properties** dialog box.

For a stripe expansion, when you add new LUNs and then set the **User Capacity** to **Current capacity**, any existing data is restriped across the original storage and any new LUNs, but you cannot use the added capacity until you change the capacity using the MetaLUN properties dialog box.

- ◆ **Maximum capacity** is the sum of the user capacity of the original LUN, metaLUN, and any added LUNs. When you select **Maximum Capacity**, the metaLUN expands to this capacity.
- ◆ **Other capacity** is a user capacity that you specify. This capacity must be equal to or greater than the existing user capacity of the LUN or metaLUN you are expanding, and equal to or less than the maximum capacity of the newly expanded metaLUN.

For example, if the metaLUN you are expanding has a user capacity of 4 Gbytes, and you are expanding this metaLUN to a maximum capacity of 7 Gbytes, in **Other capacity**, you can enter a value that is greater than or equal to 4 Gbytes and less than or equal to 7 Gbytes. The metaLUN is expanded to the capacity value that you enter.

Specifying metaLUN settings

This **Specify MetaLUN Settings** screen is available only when you first create a metaLUN. If later, you need to modify the metaLUN settings, use the **MetaLUN Properties** dialog box.

The **Specify MetaLUN Settings** screen lets you set the following settings for a new metaLUN or accept the defaults:

- ◆ **Name** — Lets you assign a name to the metaLUN.
- ◆ **Default Owner** — SP that assumes ownership of the metaLUN when the storage system is powered up (defaults to the setting of the base LUN). Select **SP A** or **SP B**.
- ◆ **Enable Auto Assign** — Enables or disables auto assign for the metaLUN. Auto assign works the same for metaLUNs as it does for LUNs. See "Enable auto assign" on page 8-8. Auto assign defaults to the setting of the base LUN.
- ◆ **Expansion Rate** — Specifies the relative importance of making additional LUN capacity available to the server. It determines the amount of resources the SP devotes to restriping any existing data across the original LUNs and the new LUNs instead of to normal

I/O activity. This rate is only important for a stripe expansion. Valid rebuild priorities are **ASAP** (default), **High**, **Medium**, and **Low**.

An expansion rate of **ASAP** or **High** restripes the data across the metaLUN faster than one with **Medium** or **Low** priority, but may degrade storage-system performance.

The wizard lets you set the expansion rate only when you first create a metaLUN. You can change the expansion rate later using the **MetaLUN Properties - General** tab. Changing the expansion rate does not affect the current expansion.

- ◆ **Element Size Multiplier** — Sets the stripe element size for a metaLUN component. You determine this by multiplying the stripe size of the first LUN in each component by the element size multiplier. The stripe element size for each component in the metaLUN can be different since the stripe size of the first LUN in each component can vary. Valid multiplier values are 1 through 255. The default multiplier is 4.

EMC recommends that you leave the element size multiplier value at 4.

- ◆ **Alignment Offset** — Specifies the offset for a LUN's RAID stripe alignment when the server operating system records private information at the start of the LUN. The default value is zero and this supports most server operating systems.

If you do not use the default value, enter a value that equals the number of sectors by which the user data is offset due to the file system or partition metadata.

Alignment Offset is available only when the LUN you are expanding meets one of the following conditions:

- The LUN is in or out of a Storage Group and has *no* existing data
- The LUN is not in a Storage Group, has existing data, but you do not want to preserve this data.
- The LUN does *not* already have an alignment offset.

Problems can arise when a server operating system records private information at the start of a LUN. This can interfere with the RAID stripe alignment so that when data I/O crosses the RAID stripe boundary, the storage-system performance is degraded. In order to properly use **Alignment Offset**, you must know the exact size of the private information. Improper use will degrade performance.

Viewing metaLUN information

Use this screen to review and verify the metaLUN summary information. This screen displays the following:

- ◆ A summary of the metaLUN settings, for example, metaLUN name, default owner, expansion.
- ◆ The type of expansion (stripe or concatenate).
- ◆ A list of LUNs that will be added to the original LUN or metaLUN.

If the information is correct, click **Finish** to begin the expansion operation.

Displaying and changing MetaLUN properties

Refer to the online help for more details on each tab of **MetaLUN Properties**.

The metaLUN properties determine the individual characteristics of a metaLUN. You set metaLUN properties when you first create the metaLUN.

You can change the following metaLUN properties after the metaLUN is created:

- ◆ Name (see page 10-15)
- ◆ Default Owner (see page 10-15)
- ◆ Expansion rate (see page 10-15)
- ◆ User Capacity (see page 10-14)
- ◆ Auto assign (see page 10-15)

To display and change the metaLUN properties, do the following:

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icons for the storage system and SP components that own the metaLUN.
3. Right-click the icon for the metaLUN, and then click **Properties**.
4. Click the **General** tab to view and change the current settings for the metaLUN.
5. Click the **Hosts** tab to do the following:
 - Display the name, IP Address and operating system for each server that has access to the metaLUN.
 - Display the physical and logical drive mapping to the metaLUN for the selected server.
 - Update server information.

This operation could be time consuming since it scans all SCSI devices for updated server information. Also, storage system performance may be affected until this operation completes. In the confirmation dialog box, Click **Yes** to continue, or **No** to cancel.

Monitoring expansion status

A concatenate expansion completes immediately, but a stripe expansion could take some time. There are two ways to help determine the status of a metaLUN stripe expansion — the **Storage** tree and the **MetaLUN Properties - General** tab.

- ◆ In the **Storage** tree, the state or icon for the expanding LUN or metaLUN changes from normal to transitioning until the expansion is complete.
- ◆ **Percent Expanded** in the **General** tab, tracks the progress of a current expansion by indicating the percent complete.

Destroying a metaLUN

Destroying a metaLUN unbinds all the LUNs in the metaLUN, and then destroys the metaLUN.



CAUTION

You will lose all data on these LUNs.

1. In the **Storage** tree, right-click the icon for the storage system in which the metaLUN resides.
2. Double-click the icon for the SP to which the metaLUN belongs.
3. Right-click the icon for the metaLUN you want to destroy and click **Destroy**.

Manager unbinds all LUNs in any metaLUN components and removes the metaLUN icon and metaLUN component icons from the **Storage** tree.

Using the added capacity of an expanded SnapView, MirrorView/A, MirrorView/S or SAN Copy LUN

You can expand the capacity of a LUN or metaLUN that is currently participating in a SnapView, MirrorView or SAN Copy session (see "Expanding a LUN" on page 10-5 or "Expanding a metaLUN" on page 10-9), but the server *cannot* access the added capacity until the expansion is complete, and you perform any or all of the following operations:

For a stripe expansion, any existing data is restriped across the original LUN or metaLUN and the added LUNs *during* the SnapView, MirrorView or SAN Copy session. This restriping is invisible to the server.

For LUNs in a Snapshot session

To access the added capacity of an expanded snapshot source LUN, do the following:

1. Stop *all* snapshot sessions and destroy *all* snapshots associated with the expanded source LUN (see the online help or the SnapView Administrator's Guide).
2. Use the **MetaLUN Properties - General tab** to change the capacity of the metaLUN.
3. Make sure the server operating system can utilize the added capacity.
4. You can now start new sessions and create new snapshots with the expanded LUNs (see the online help or the SnapView Administrator's Guide).

For LUNs in a clone group

You must expand the clone and the clone source LUN to the same size.

To access the added capacity of an expanded clone, or clone source LUN, do the following:

1. For the clone, remove the clone from the clone group; for a clone source LUN, destroy the clone Group.
2. Use the **MetaLUN Properties - General tab** to change the capacity of the metaLUN.
3. Make sure the server operating system can utilize the added capacity.
4. Reconstruct the clone group with the expanded LUNs and then perform a full synchronization of the LUNs in the clone group.

For LUNs in mirrors

You must expand both the primary and secondary images of a mirror to the same size.

To access the added capacity of an expanded primary and secondary MirrorView LUN, do the following:

1. Remove the expanded secondary image from the mirror, then destroy the expanded primary image and the mirror (see the online help or the MirrorView Administrator's Guide.)
2. Use the **MetaLUN Properties - General tab** to change the capacity of the metaLUN.
3. Make sure the server operating system can utilize the added capacity.
4. Reconstruct the mirror with the expanded LUNs which will require a full synchronization of the mirror.

For LUNs in SAN Copy sessions

If you expand the SAN Copy source logical unit, make sure that the destination LUNs are of equal or greater capacity.

To access the added capacity of an expanded SAN Copy source LUN or destination LUN, do the following:

1. For any sessions that the expanded LUNs are participating in, wait until the SAN Copy sessions complete, or stop the sessions, and then remove these sessions. (see the online help, or the SAN Copy Administrator's Guide).
2. Use the **MetaLUN Properties - General tab** to change the capacity of the metaLUN.
3. Make sure the server operating system can utilize the added capacity.
4. Recreate the SAN Copy sessions that you removed and then start the sessions. (see the online help, or the SAN Copy Administrator's Guide).

Managing the Reserved LUN Pool for LUN Replication

This chapter explains how to configure the reserved LUN pool, which is required for running replication software such as SnapView and SAN Copy.

Major sections are

- ◆ Introduction to the reserved LUN pool 11-2
- ◆ Estimating the reserved LUN pool size 11-3
- ◆ Configuring the reserved LUN pool 11-5
- ◆ Monitoring reserved LUN pool usage 11-7
- ◆ Displaying reserved LUN pool properties 11-8

Introduction to the reserved LUN pool

The reserved LUN pool works with replication software, such as SnapView, SAN Copy and MirrorView/A, to store data or information required to complete a replication task. The reserved LUN pool consists of one or more private LUNs, and each SP has its own reserved LUN pool. Since the LUNs in the reserved LUN pool are private LUNs, they cannot belong to Storage Groups and a server cannot perform I/O to them.

Before starting a replication task, the reserved LUN pool must contain at least one LUN for each source LUN that will participate in the task. You can add any LUNs that are available to either SP's reserved LUN pool. Each SP manages its own LUN pool space and assigns a separate reserved LUN (or multiple LUNs) to each source LUN. All replication software that uses the reserved LUN pool shares the resources of the reserved LUN pool. For example, if you are running an incremental SAN Copy session on a LUN that belongs to SP A and a SnapView session on another LUN that belongs to the same SP, the reserved LUN pool for SP A must contain at least two LUNs - one for each source LUN. If both sessions are running on the same source LUN, the sessions will share a reserved LUN.

Allocation of reserved LUNs

Each SP manages its own reserved LUN pool space and allocates reserved LUNs on a per-source-LUN basis, not a per-session basis. For example, if two sessions are running on a source LUN, the SP assigns reserved LUNs, as needed, to the source LUN, and both sessions share these LUNs.

The assignment of reserved LUNs is based upon the first available free LUN in the SP's LUN pool. An adequate number of reserved LUNs is essential because the software will terminate sessions if no free LUNs are available in the SP's LUN pool. If you are running multiple sessions per source LUN and your reserved LUN fills up (runs out of space) and no free LUNs are available in the SP's LUN pool, the session that is trying to allocate a reserved LUN from this pool will terminate.

How SPs manage the reserved LUN pool

When you start the first session on a source LUN, reserved LUNs from the LUN pool of the SP that currently owns the source LUN are assigned to the source LUN. All subsequent sessions started on this source LUN will use reserved LUNs from the same SP's LUN pool, even if the source LUN trespasses to the peer SP (unless you stop all sessions on the trespassed source LUN).

For example, if you start the first session on a source LUN that belongs to SP A, and the source LUN is trespassed to SP B, the reserved LUNs assigned to this source LUN also trespass to SP B. If you start a second session on the trespassed source LUN (which now belongs to SP B), the second session will still use the same reserved LUNs that were trespassed from SP A. If the sessions running on the trespassed source LUN need additional reserved LUNs, reserved LUNs from SP A's LUN pool are used, not SP B's. If, however, you stop all sessions on the trespassed source LUN, and then start a new session on this LUN, the new session uses reserved LUNs from SP B's LUN pool, not SP A's.

Estimating the reserved LUN pool size

Each reserved LUN can vary in size. However, using the same size for each LUN in the pool is easier to manage because the reserved LUNs are assigned without regard to size; that is, it assigns the first available free LUN in the SP's reserved LUN pool. Since you cannot control which reserved LUNs it uses for a particular session, EMC recommends that you use more, smaller reserved LUNs. By doing this, Manager is better able to allocate the reserved LUN space. If a session requires additional reserved LUN space, Manager dynamically assigns a free reserved LUN from the SP's LUN pool.

Before binding LUNs for the reserved LUN pool, think about the size of the LUNs that will use LUNs in the reserved LUN pool. This will help you estimate a suitable size for the reserved LUN pool LUNs.

Estimating a suitable reserved LUN pool size

To help determine the size of the reserved LUN pool and the reserved LUNs, do the following:

1. Add the total size of source LUNs in SP A or SP B.
2. Multiply that number by 10%. This value is the reserved LUN pool usage.
3. Multiply the reserved LUN pool usage amount by two (this is a buffer for any unanticipated I/O to the source LUN(s)).

The result is a reasonable estimate for the size of the reserved LUN pool.

4. To determine the size of each reserved LUN, divide the reserved LUN pool size by two times the number of source LUNs. The result is the size of each reserved LUN.

To protect your sessions Snapview, SAN Copy or MirrorView /A) from being terminated (if no free reserved LUNs were available in the pool), the number of reserved LUNs in the pool should always be two times the number of source LUNs.

In addition to using the size of the source LUN to determine a suitable size for the reserved LUN pool, you may want to consider the duration time for the replication session in your calculations. A longer duration time could increase the need for more reserved LUN pool space since it is possible that more areas will be written to over a longer period of time.

Configuring the reserved LUN pool

You must configure the reserved LUN pool before you start replication tasks such as SnapView sessions or incremental SAN Copy sessions. Each SP has its own reserved LUN pool and you can configure or change the LUN pool size of an SP's LUN pool by adding or removing LUNs. If you plan to start sessions for LUNs owned by SP A and LUNs owned by SP B, configure the reserved LUN pool for each SP.

The software allocates the reserved LUNs on a per-source LUN basis, not a per-session basis.

To configure the reserved LUN pool, do the following:

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icon for the storage system for which you want to configure the reserved LUN pool, and then double-click the **Reserved LUN Pool** icon.
3. Right-click either the **SP A** or **SP B LUN Pool** icon, and then click **Configure**.

The **Configure Reserved LUN Pool** dialog box opens.

In **Chunk Size (Blocks)**, the reserved LUN chunk size (size of each write) for both SPs is set to the default size of 64K (128 blocks in Navisphere). You cannot change this value.

4. In **Available LUNs**, select the LUNs that you want to add to the reserved LUN pool. **Available LUNs** lists only those LUNs that are eligible for inclusion in the reserved LUN pool. If there are no LUNs in the list, or if none of the available LUNs are the necessary size, you must bind new LUNs, and then add them to the reserved LUN pool (see *"Creating LUNs on RAID groups" on page 8-16*).
5. To add the selected LUNs to SP A's reserved LUN pool, click **Add to SP A LUN Pool**; to add LUNs to SP B's reserved LUN pool, click **Add to SP B LUN Pool**.

The selected LUNs move to the **SP LUN Pool LUNs** list for SP A or SP B and the value in **Modified Capacity** is updated to reflect the changes.

6. To remove LUNs from the **SP LUN Pool LUNs** list, select the LUNs you want to remove, and then click **Remove LUN**.

You can remove only free (unallocated) reserved LUNs from the **SP LUN Pool LUNs** list.

7. When you have added an appropriate number of LUNs to the reserved LUN pool, click **OK** to apply the changes and close the dialog box.

Manager displays an icon for each reserved LUN in the **Storage** tree under **Reserved LUN Pool > SP A** or **SP B LUN Pool**.

Monitoring reserved LUN pool usage

To monitor the reserved LUN pool usage and to determine if it is the necessary size, do the following:

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icon for the storage system on which the reserved LUN pool resides, and then double-click the **Reserved LUN Pool** icon.
3. Right-click the icon for the SP of the reserved LUN pool you want to monitor, and then click **Properties**.

The **Reserved LUN Pool Properties** dialog box opens.

4. Under **Allocated LUN Pool**, check **LUN Pool Usage**.

If the SP usage registers at 80% to 90%, you may want to increase the size of the reserved LUN pool.

5. Click **Help** for information on the properties in the dialog box, or click **Close** to close the dialog box.

If you determine that the reserved LUN pool is running out of space, you can do one of following:

- ◆ stop a session to free its LUN space, or
- ◆ add LUNs to the reserved LUN pool while the session is active

If the reserved LUN fills up, the software automatically terminates the session that encountered the error, logs an error, and releases the reserved LUN space that session used. If you are running multiple sessions per source LUN(s), you must stop all sessions on that source LUN(s) in order to free reserved LUN space for use by sessions on a different source LUN(s).

You can configure Navisphere Manager to send an event notification when the reserved LUN pool becomes 50, 75, 80, 85, 90, or 95 percent full. The application will generate the following warning event code and description when each of the reserved LUN pool capacities is reached:

Event code	Description
0x71004000	The reserved LUNs assigned to LUN X have become Y% percent full. where X is the source LUN and Y is either 50, 75, 80, 85, 90, or 95.

Displaying reserved LUN pool properties

All reserved LUN pool properties are read only. If you want to change the size of an SP's LUN pool, refer to "Configuring the reserved LUN pool" on page 11-5.

To display the properties of the reserved LUN pool for SP A and SP B, do the following:

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icon for the storage system for which you want to display the reserved LUN pool properties, and then double-click the **Reserved LUN Pool** icon.
3. Right-click either the **SP A** or **SP B LUN Pool** icon, and then click **Properties**. The **Reserved LUN Pool Properties** dialog box opens and the following information is available:

Allocated LUN Pool

- Source LUN
- LUN Pool LUNs
- LUN Pool Size
- LUN Pool Usage

Free LUN Pool

- Name
- Size
- Drive Type

4. Click **Close** to close the dialog box.

Click **Help** for information on the properties in the dialog box.

Setting Up and Using Event Monitor

This chapter introduces Event Monitor, which allows you to track the status of storage systems connected to servers on a direct attach storage (DAS) or storage area network (SAN) configuration. It also describes the Event Monitor Wizard and Event Viewer.

For security reasons, we recommend that you place all storage systems running the Storage Management Server software in a domain before you configure them.

To configure the storage system you must be logged in as global or local administrator or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes the following topics:

- ◆ Event Monitor overview 12-2
- ◆ Event Monitor terminology 12-4
- ◆ Distributed and centralized monitoring 12-5
- ◆ Event Monitor configuration overview 12-10
- ◆ Configuring Event Monitor using the wizard 12-11
- ◆ Configuring Event Monitor using menu options 12-17
- ◆ Testing responses 12-27
- ◆ Testing configurations 12-28
- ◆ Deleting customized responses 12-29
- ◆ Importing and exporting templates 12-29
- ◆ Viewing and modifying template properties 12-31
- ◆ Deleting templates 12-32
- ◆ Applying templates to storage systems 12-33
- ◆ Using Event Monitor 12-34

Event Monitor overview

The Event Monitor is an enterprise tool that supports Centralized or Distributed monitoring of storage systems in a heterogeneous environment. The Event Monitor software consists of two distinct parts: the Event Monitor Graphical User Interface (GUI), and the Event Monitor.

The Event Monitor Graphical User Interface (GUI) is part of Navisphere Manager and runs on the web browser. The user interface provides you with an intuitive tool to setup responses for events and to choose which storage systems to observe. The user interface lets you customize a configuration to use any of the supported notification methods. It can be easily configured to e-mail, page, or send an SNMP trap to an industry standard event-management tool. The user interface need only be used when setting up configurations or viewing the Event History log.

Event Monitor resides on both the Navisphere SP Agent and Host Agent, and is available on many operating systems. Once configured, the Event Monitor runs continuously as a service or daemon, observing the state of all specified storage systems and notifying you when selected events occur.

Event Monitor tracks the status of storage systems connected to servers on a direct attach storage (DAS) or storage area network (SAN). You can specify the storage systems for which Event Monitor logs status and the range of event messages it reports. You can also configure Event Monitor to notify you when selected events occur using any combination of the following methods:

- ◆ E-mail
- ◆ Paging via modem or E-Mail
- ◆ Sending an SNMP trap to an industry-standard network-management tool such as HP OpenView™, Tivoli® Netview®, and CA-Unicenter® TNG.
- ◆ Custom response

The types of responses that Event Monitor allows vary between server operating systems and on whether you are using centralized or distributed monitoring. With distributed monitoring, SP Agents monitor only that storage system in which the SP Agent resides, or Host Agents monitor all attached storage systems. With centralized monitoring, a monitoring Agent monitors storage systems from a central location. For both configurations, the monitoring Agent notifies the client of events which occur on the monitored storage systems.

Table 12-1 defines the responses available from the various platforms in server-resident centralized monitoring environments.

Table 12-1 Centralized monitoring response matrix

Platform	LogToSysLog	E-mail	PageViaEmail	SNMP	PageViaModem	Custom
Windows	X	X	X	X	X	X
Solaris	X	X	X	X		X
AIX	X	X	X	X		X
LINUX	X					X
IRIX	X					
NetWare	X					
HP-UX	X	X	X	X		X

Table 12-2 defines the responses available from the various platforms in server-resident distributed monitoring environments.

Table 12-2 Distributed monitoring response matrix - host resident

Platform	LogToSysLog	E-mail	PageViaEmail	SNMP	PageViaModem	Custom
Windows	X	X	X	X	X	X
Solaris	X	X	X	X		X
AIX	X	X	X	X		X
LINUX	X					X
IRIX	X					
NetWare	X					
HP-UX	X	X	X	X		X

Table 12-3 defines the responses available from SP-resident distributed monitoring environments.

Table 12-3 Distributed monitoring response matrix - SP resident

Storage System Type	LogToSysLog	E-mail	PageViaEmail	SNMP	PageViaModem	Custom
CX-Series and FC4700-Series		X	X	X		

Event Monitor terminology

Server — A host that uses one or more storage systems for data storage retrieval.

Agent — A server-based or SP-based daemon application that communicates between Navisphere Server applications and storage systems. The Agent is the only application that communicates directly with the storage system.

Centralized Monitoring — The process where a single monitoring Agent monitors several storage systems (via Host Agents or SP Agents) and notifies you when selected events occur.

Distributed Monitoring — The process where each Host Agent or SP Agent monitors its own storage system and notifies you when selected events occur.

Event Monitor — The part of the Agent program that can monitor itself or other servers.

Managed Storage System — A storage system that EMC software manages.

Monitor Station — A server with Event Monitor configured as a centralized monitor.

Monitored Server — A server whose storage-system status is monitored.

Self-monitoring Agent — Agent that monitors the status of the storage system on which it resides (SP Agent), or the storage system to which it is connected (Host Agent), and then notifies you of selected events for each monitored storage system. Self-Monitoring Agents are used in distributed monitoring applications.

Centralized Monitoring Agent — Host Agent that monitors storage systems on the network by communicating with other Host Agents or SP Agents. The Centralized Monitoring Agent notifies you when events occur on any of the storage systems it is monitoring.

Monitored Agent — Agent that monitors the status of the storage system on which it resides (SP Agent), or the storage system to which it is connected (Host Agent), and reports events to the Centralized Monitoring Agent.

Distributed and centralized monitoring

You can configure a system to use either distributed or centralized monitoring. In a distributed monitoring environment, each SP Agent or Host Agent monitors its own storage system or the storage systems to which it is connected, and notifies you when selected events occur (refer to Figure 12-1). In a centralized monitoring environment, a single monitor station monitors several storage systems (via Host Agents or SP Agents) and then notifies you when selected events occur (refer to Figure 12-2).

Monitoring guidelines

We recommend that you follow these guidelines when configuring your monitoring environments:

- ◆ All storage systems with pre-6.1 SP Agents and all storage systems without Storage Management Server software installed should be centrally monitored.

The revision of the Central Agent must be greater than or equal to the revision of the monitored servers.

- ◆ Storage systems with 6.1 and higher SP Agents can be monitored using centralized or distributed monitoring.
- ◆ Within a domain, all monitoring should be either distributed or centralized. Do not mix the two types of monitoring within a domain.


Configuring distributed monitoring

To use Navisphere Manager 6.X to configure distributed monitoring, do the following:

1. Decide which Agents you want to configure to use distributed monitoring.
2. Use the **Configuration Wizard** or the right-click menu options to create and apply templates to the selected Agents.

The templates you create and apply to the Agents specify the storage system events that you want to monitor and the method of event notification. You can view the Event Logs, Response Logs and Message Files for each of the monitored Agents using the Event Monitor GUI. Depending on the operating system, you must create one or more templates for each of the monitored Agents.

In Figure 12-1, each Agent is monitoring the storage systems attached to it (the Windows NT and UNIX servers), or monitoring the storage system in which it resides (the Storage Management Server), and then notifies you when selected events occur.



EMC2399

Figure 12-1 Distributed monitoring system

Configuring centralized monitoring for a single domain

To use Navisphere Manager 6.X to configure centralized monitoring for a single domain, do the following:

1. Install the Host Agent on a host that is not managing any storage systems. This becomes the Centralized Monitoring Agent.

You cannot use an SP Agent as the Centralized Monitoring Agent.

2. Create a domain, making sure to include all the storage systems running Storage Management Server software that you want to monitor.
3. Use the **Portal Configuration** dialog box to
 - add a portal (use one of the systems in the domain).
 - assign the Centralized Monitoring Agent to the portal configuration using the **Add Storage System** right-click option.
 - add, to the portal, any Agents managing legacy systems that you want to monitor.

For information on configuring portals, refer to Chapter 17, "Managing Legacy Storage Systems and NAS Systems".

4. When you finish configuring the portal, click **OK** in the **Portal Configuration** dialog box.

Manager places an icon for the monitoring Agent in the **Monitors** tab of the **Enterprise Storage** window, and adds icons for the legacy systems to the **Storage** tree.

5. In order to issue commands to the centralized monitoring Agent, you must add the IP address(es) of the portal to the monitoring Agent's agent configuration file. Refer to "Adding IP addresses to the agent configuration files" on page 17-8.
6. You can now use the **Configuration Wizard** or the right-click menu options to add monitored Agents to the centralized monitoring configuration, and create and apply Centralized Event Monitoring templates to the monitoring agent.

Do *not* add another Centralized Monitoring Agent to the monitoring configuration.

In Figure 12-2, the Host Agent in the Monitor Station serves as the Centralized Monitoring Agent. EMC recommends that you do *not* use a server that is managing storage systems as the Centralized Monitoring Agent (the UNIX® or Solaris™ servers).




Figure 12-2 Centralized monitoring system

Configuring centralized monitoring for multiple domains

You must use the same Centralized Monitoring Agent to monitor storage systems from multiple domains. This monitoring Agent resides on a host that is not managing any storage systems.

1. Install the Host Agent on a host that is not managing any storage systems. This becomes the Centralized Monitoring Agent.

You cannot use an SP Agent as the Centralized Monitoring Agent.

2. For *each* domain with storage systems that you want to monitor, use the **Portal Configuration** dialog box to
 - add a portal (use one of the systems in the domain).
 - assign the Centralized Monitoring Agent to the portal configuration using the **Add Storage System** right-click option.
 - add, to the portal, any Agents managing legacy systems that you want to monitor.

For information on configuring portals, refer to Chapter 17, "Managing Legacy Storage Systems and NAS Systems".

3. When you finish configuring each portal, click **OK** in the **Portal Configuration** dialog box.

Manager places an icon for the monitoring Agent in the **Monitors** tab of the **Enterprise Storage** window, and adds icons for the legacy systems to the **Storage** tree.

4. In order to issue commands to the monitoring Agent, you must add the IP address(es) of each portal to the monitoring Agent's agent configuration file. Refer to "Adding IP addresses to the agent configuration files" on page 17-8.
5. You can use the **Configuration Wizard** or the right-click menu options to add monitored Agents to the centralized monitoring configuration, and create and apply centralized Event monitoring templates to the monitoring agent. To add monitored Agents in other domains, refer to "Adding Agents outside the domain" on page 12-17.

Do *not* add another Centralized Monitoring Agent to the monitoring configuration.

Event Monitor configuration overview

You can use the **Configuration Wizard** or the right-click menu options to set up Event Monitor configurations. The Wizard uses a series of intuitive screens to help you set up distributed or centralized monitoring environments. It lets you select monitoring Agents and then create and assign templates to these Agents. A template contains events, responses, and message formatting and can be mapped to one or more storage systems. This template specifies the storage system events that you want to monitor and the method of event notification. The Wizard is an ideal tool to use for setting up new Event Monitor configurations. Refer to "Configuring Event Monitor using the wizard" on page 12-11.

If you configure Event Monitor to use multiple templates to respond to events from a storage system it is monitoring, Event Monitor may generate multiple responses for a single event.

In addition to setting up new Event Monitor configurations, the right-click menu options let you view and modify any existing configurations. Refer to "Configuring Event Monitor using menu options" on page 12-17.

Configuring Event Monitor using the wizard

Use the Event Monitor Wizard to set up new Event Monitor configurations.

The Wizard automatically creates and applies the correct version of template to a centralized monitoring Agent. For distributed monitoring, you must select the version of the Agent to which you want to apply the templates.

You can apply a version 6.2 or higher template only to a 6.2 or higher Agent and a version 6.1 or earlier template to a 6.1 or earlier Agent. For domains containing both 6.2 and 6.1 Agents, you must run the wizard once to set up the 6.1 Agents and then again to set up the 6.2 Agents.

Starting the wizard

1. On the **Tools** menu, select **Event Monitor > Configuration Wizard**.

or

In the **Enterprise Storage** dialog box, click the **Monitors** tab, right-click **Templates**, and select **Configuration Wizard**.

2. Read the setup steps in the **Start Wizard** screen and then click **Next** to continue.

Selecting the monitoring type

3. In the **Select Event Monitor Type** screen enter the template name, and select the type of monitoring you want: distributed or centralized.
4. Click **Next** to continue.
 - If you selected **Distributed Monitoring**, the **Template Option** screen opens.
 - If you selected **Centralized Monitoring**, the **Centralized Monitoring** screen opens.

For distributed monitoring

5. For **Distributed Monitoring**, do the following:
 - a. In the **Template Option** screen, select Navisphere 6.1 or earlier if you will be applying this template to a version 6.1 or earlier Agent, or select Navisphere 6.2 or greater if you will be applying this template to a version 6.2 or higher Agent.
 - b. Click **Next** to select the storage systems that you want to monitor.

Event Monitor displays only those storage systems with an SP Agent that matches the version of the Agent that you selected.

For centralized monitoring

- c. Click **Next** to open the **Select Events by Category** screen, and then skip to step 7.

6. For Centralized Monitoring, do the following:

The wizard automatically creates and applies the correct version of template to a centralized monitoring Agent.

- a. Select the host that you want to be the Monitor Station (perform the central monitoring), and then click **Next**.
- b. In the **Select Storage System** screen, select the storage systems that you want to monitor.
- c. Click **Monitor Options**, to do any of the following in the **Event Monitor Agent Properties** screen:
 - change the timeout minutes for Centralized Monitoring.
 - change the maximum log file size for stored events. (maximum is 25 Mbytes).
 - For configuring multiple domains only - add Agents from other domains that you want to be monitored using the same Centralized Monitoring Agent. Refer to "Configuring centralized monitoring for multiple domains" on page 12-9.
- d. Click **OK** to close the **Event Monitor Agent Properties** screen and return to the **Select Storage System** screen.
- e. Click **Next** to open the **Select Events by Category** screen.

Selecting events to monitor

7. In the **Select Events by Category** screen, select one or more event categories that you want to trigger responses, and then click **Next**.
8. In the **Select Events by Severity** screen, specify the severity of the errors you want reported by either selecting **General** or **Explicit Events**.

When **Explicit Events** is selected, all options in the **General** field are dimmed and unavailable.

- a. If you select **General**, do the following:
 - Select one of the four severity code options in the **General** field:
 - Only the most critical error conditions**
 - All error conditions**
 - All error conditions and warnings**
 - All errors, warnings and information messages**
 - If you are creating a template for a version 6.1 or earlier Agent, **Advanced** is dimmed and unavailable.
 - If you are creating a template for a version 6.2 or higher Agent, click **Advanced** to open the **General Event Threshold Advanced Selection** screen to specify the threshold count, threshold interval and response delay for event notification.

Threshold count is the number of times an event must occur within a specified amount of time in order for Event Monitor to generate a response. The default threshold count value is 1; the minimum value is 1, and the maximum value is 50.

Threshold interval is the time interval over which an event must occur for a specified number of times in order for Event Monitor to generate a response. Depending on the version of Agent, this interval is measured in minutes.

Response delay is the time that Event Monitor waits to send a repeat notification for an event after sending the initial notification.
 - When you are finished, click **OK**.
 - b. If you select **Use Explicit Events**, click **Advanced** to open the **Advanced Event Code Selection** screen.
 - Select the filter by which you want to explicitly select events, expand the event code that you want to select, and select the check boxes for the events that you want to be monitored.
 - When you are finished, click **OK**.
9. Click **Next** to open the **Select Response** screen.

Selecting the response method

10. In the **Select Response** screen, select the response methods that you want the system to use to notify you about system events: **Send E-mail**, **Paging Service**, or **Send SNMP Trap**, and then click **Next** to continue.

You can specify any of these responses in any combination. For more information about responses, refer to Chapter 6, "Configuring the Remote Host Agent".

11. Complete the information for each response you want.

E-mail response - Go to the next section, "E-mail responses" on page 12-14.

Pager response - Go to "Pager responses" on page 12-14.

Pager services - Go to "Paging service using e-mail" on page 12-15.

SNMP response - Go to "SNMP response" on page 12-16.

E-mail responses

12. In the **E-Mail** screen, do the following:

- a. Enter the appropriate information in **E-Mail To**, **Cc**, **From**, **Subject**, and **SMTP Server**.

If you want to monitor a CX-Series or FC4700-Series storage system using this template, enter the IP address of the mail server for this system in **SMTP Server**.

- b. Do one of the following:

- Select the **Use Default Message** check box to specify that Event Monitor E-mail the message using the default message text and format.
- Clear the **Use Default Message** check box to enable the message box, and enter a customized message.

- c. Click **Next** to set up the next response method or to proceed to the **Summary** screen in step 16.

Pager responses

13. In the **Paging** screen, do the following:

- a. Select **Use a modem to call the paging service**.
- b. Enter the appropriate information in the **Phone #** and **Message #** fields.

- c. Click **More** to set up the advanced paging features: **COM Port, Message Delay, and Dial Command**, and **Init Command**.

Before you set up the advanced paging features, you must know the port to which your modem is configured.

For more information about Init commands, go to www.modemhelp.org.

- d. Click **OK** to set up the advanced settings and close the **Paging Advanced Set Up** screen.
- e. Click **Next** to set up the next response method, or to proceed to the **Summary** screen in step 16.

Paging service using e-mail

14. In the **Paging** screen, do the following:

- a. Select **Use E-Mail to contact the paging service**.
- b. Enter the appropriate information in **E-Mail To, Cc, From, Subject**, and **SMTP Server**.

If you want to monitor a CX-Series or FC4700-Series storage system using this template, enter the IP address of the mail server for this system in **SMTP Server**.

- c. Do one of the following:
 - Select the **Use Default Message** check box to specify that Event Monitor E-mail the message using the default message text and format.
 - Clear the **Use Default Message** check box to enable the message box, and enter a customized message. To clear the default message box, select the text and drag down four lines, and press **Delete**.
- d. Click **Next** to set up the next response method, or to proceed to the **Summary** screen in step 16.

SNMP response

15. In the SNMP screen, do the following:
 - a. Enter the appropriate information in the **SNMP Management Host** and **Community** fields.
 - b. Click **Next** to complete the SNMP response setup and proceed to the **Summary** screen.

Refer to "Receiving SNMP traps" on page 12-25 for information on severity levels in SNMP traps.

16. The wizard opens the **Summary** screen, which lets you review the servers, event types, event severities, and responses you defined during this session.
17. Review the information displayed in this screen. If it is satisfactory, then click **Finish** to create the template file.

Event Monitor will apply the templates to all the selected event monitoring Agents, that will monitor the storage systems.

If you want to make changes, click **Back** to return to the screen that you want to change; make your changes, and then click **Next** until you return to the **Summary** screen.

Configuring Event Monitor using menu options

In addition to setting up new Event Monitor configurations, such as adding monitored Agents, creating templates and setting up event notification responses, you can also use the right-click menu options to view and modify existing configurations.

Adding monitored Agents

Event Monitor lets you add monitored Agents that are within or outside the domain, to a centralized monitoring configuration.

You must be a privileged user to add Agents. Do not add the local host to this dialog box.

Adding Agents within the domain

The name or IP address of hosts within the domain, appear in a list format.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. In the **Monitors** tree, right-click the icon for the monitoring Agent and select **Monitor Options** to open the **Event Monitor Agent Properties** dialog box.
3. From **Available Systems**, select the host you want to add.
4. Click the **Right Arrow**, located between **Available Systems** and **Selected Systems**.

The host now appears in **Selected Systems**.

Adding Agents outside the domain

You type the name or IP address of a host, to add an Agent outside the domain.

You add Agents in other domains if you are configuring centralized monitoring for multiple domains. Refer to "Configuring centralized monitoring for multiple domains" on page 12-9.

1. If necessary, repeat steps 1 and 2 above.
2. In **Name** or **IP Address**, enter the name or IP address of the host you want to add.
3. Click the **Right Arrow**, located to the right of **Name** and **IP Address**.

The host now appears in **Selected Systems**.

Removing monitored Agents

When you no longer want a centralized monitoring Agent to monitor specific Agents, you can remove the Agents using the **Event Monitor Agent Properties** dialog box.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. In the **Monitors** tree, right-click the icon for the monitoring Agent and select **Monitor Options**.
3. In the **Event Monitor Agent Properties** dialog box, select the Agent that you want to remove from **Selected Systems**.
4. Click the **Left Arrow**, located between **Available Systems** and **Selected Systems**.
5. Click **OK**.

Configuring log file size and network timeout

The **Monitor Options** dialog box lets you set the following properties for the Monitoring Agent:

- ◆ Log file size where each record is 512 bytes or 0.5 Kbytes. The log file is circular and automatically overwrites the oldest entries. The default log size is 25 Mbytes or 50,000 records.
- ◆ Network timeout for centralized monitoring. This timeout occurs when the Monitoring Agent sends a request to one of the Agents it is monitoring, and the Agent does not respond within a specified period of time.

Do not confuse the network timeout for centralized monitoring with the network timeout for Manager that you must set in Manager. A network timeout for Manager occurs when Manager sends a request to an Agent, and the Agent does not respond within a specified period of time.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. In the **Monitors** tree right-click the monitoring host icon and select **Monitor Options**.
3. Select a time for **Network Timeout for Central Monitoring**.
4. Select the log file size.
5. Click **OK**.

The existing log file is backed up and cleared when the log file is modified.

Creating new templates without the wizard

A template contains events, responses, and message formatting. A template can be mapped to one or more storage systems. A custom template may contain additional statements that describe various data elements needed by a custom response.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click **Templates**.
3. Select **Create New Template** or, on the **Tools** menu, select **Event Monitor > Create New Template** to open the **Option to Create Template** dialog box.

This dialog box lets you select the version of the Agent to which you will be applying the new template

4. Click **OK** to open the **Template Properties - General** tab dialog box
5. On the **General** tab, type the name of the template in the **Template Name** field.

The length of a template filename is limited to 38 characters. In addition, the name must not start with a number, and it must contain characters in the following range: [a...z], [A...Z], [0...9]. If you import an existing template with a filename of illegal length or containing non-supported characters, the file is not replicated across the domain.

6. Under **Events**, select either **General** or **Explicit Events**.

When **Explicit Events** is selected, all options in the **Event Severity** and **Event Category** fields are dimmed and unavailable.

7. If you have selected **General Events**
 - a. Under **Event Severity**, specify the severity of the errors you want reported.
 - b. Under **Event Category**, select the categories of events that you want this template to monitor.

- c. If you are creating a template for a version 6.2 or higher Agent, click **Advanced** to open the **General Event Threshold Advanced Selection** dialog box and then specify the threshold count, threshold interval and response delay for event notification.
 - **Threshold count** is the number of times an event must occur within a specified amount of time in order for Event Monitor to generate a response. The default threshold count value is 1; the minimum value is 1, and the maximum value is 50.
 - **Threshold interval** is the time interval over which an event must occur a specified number of times in order for Event Monitor to generate a response. Depending on the version of Agent, this interval is measured in minutes.
 - **Response delay** is the time that Event Monitor waits to send repeat notification for an event after sending the initial notification.
8. If you have selected **Explicit Events**, click **Advanced**. The **Advanced Event Code Selection** dialog box opens.
 - a. Select the filter by which you want to explicitly select events, expand the event code that you want to select, and select the check boxes for the events that you want to be monitored.
 - b. Click **OK**.
9. In the **Action for events** field, select the action that you want the system to use to notify you of the events.
 - **Log to System Log**—Event Monitor logs all the events to the system log.

 Windows NT Agents log events to the System Event log.
 - **Single Notification for Multiple Events**—Event Monitor notifies you only once for multiple events. For example, if you set up E-mail as the event type, Event Monitor notifies you once via E-mail of all the events that have occurred in the poll cycle. Event descriptions will be contained in the body of the message.

 If you select this check box, Event Monitor cannot send SNMP traps.

10. Click the **Add Response** button if the existing responses do not meet your needs.

You can create and add customized responses which let you configure Event Monitor to execute a program. Event Monitor adds the customized response tab next to the other response tabs.

Setting up responses for the templates

This section describes how to set up E-mail, pager, and SNMP responses to the events you are monitoring. In addition to letting you configure the response you want to any events, each response tab includes a **Test** button. Refer to "Testing responses" on page 12-27.

You must enter a DNS or host file entry for each hostname you enter in a response tab. You can use IP addresses for hostname entries.

Setting up e-mail responses

If you set up E-mail as a response, Event Monitor notifies you of system events by SMTP E-mail. An SMTP server must be configured to receive the E-mail response.

To set up the E-mail response:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template for which you want to set up the E-mail response.
3. Select **Properties**.
4. In the **Template Properties** dialog box, click the **E-Mail** tab.
5. Enter the appropriate information in the **E-Mail To**, **Cc**, **From**, **Subject**, and **SMTP Server** fields.

If you want to monitor a CX-Series or FC4700-Series storage system using this template, enter the IP address of the mail server for this system in **SMTP Server**.

6. Do one of the following:
 - Select the **Use Default Message** check box to specify that Event Monitor E-mail the message using the default message text and format.

By default, the **Use Default Message** box is checked.

 - Clear the **Use Default Message** check box to enable the message box, and enter a customized message.
7. Click **Test** to test your response setup. For more information on testing responses, refer to the section "Testing responses" on page 12-27.
8. In the **Select Test Host** dialog box, select the Agent you want to use from the drop-down menu.
9. Click **OK** to apply the Agent and close the **Select Test Host** dialog box.
10. Click **OK** to close the **Template** dialog box.

Setting up paging responses

Event Monitor lets you set up Paging responses via modem or E-mail. If you set up paging as a response, Event Monitor notifies you of system events by pager.

If you are setting up paging responses by modem, be sure your modem is set up properly. If you are setting up paging responses by E-mail, be sure you have a pager service server that can send E-mail messages to the pager.

Setting up paging responses by modem

Follow these steps to set up the paging response via modem:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template for which you want to set up paging responses.
3. Select **Template Properties**.
4. Click the **Paging** tab.
5. Select **Use Modem to call the paging service**.
6. Enter the appropriate information in the **Phone #** and **Message #** fields.

7. Click **More** to set up the advanced paging features. For more information, refer to the next section "Setting up advanced paging features".
8. Click **Test** to test the pager response. For more information, refer to the section "Testing responses" on page 12-27.
9. In the **Select Test Host** dialog box, select the Agent with dial-up modem that you want to use for testing from the drop-down menu.
10. Click **OK** to close the **Select Test Host** dialog box.
11. Click **OK** to close the **Template** dialog box.

Setting up advanced paging features

Before you set up the advanced paging features, you must know the port to which your modem is configured. This modem must be different from the EMC Phone Home modem.

To set up advanced paging features:

1. In the **Template Properties** dialog box, click the **Paging** tab.
2. Click **More**.

The **Paging Advanced Set Up** dialog box opens.

3. Use the drop-down menus to select the appropriate information in the **COM Port**, **Message Delay**, and **Dial Command** fields.
4. Enter the Init string in the **Init Command** field. For more information about Init commands, go to www.modemhelp.org.
5. Click **OK** to set up the advanced settings and close the **Paging Advanced Set Up** dialog box.

Setting up pager responses by e-mail

1. Follow steps 1 - 4 shown in "Setting up paging responses by modem" on page 12-22.
2. Select **Use E-Mail to contact the paging service**.
3. Enter the appropriate information in **E-Mail To**, **Cc**, **From**, **Subject**, and **SMTP Server**.

If you want to monitor a CX-Series or FC4700-Series storage system using this template, enter the IP address of the mail server for this system in **SMTP Server**.

4. Do one of the following:
 - Select the **Use Default Message** check box to specify that Event Monitor E-mail the message using the default message text and format.
 - Clear the **Use Default Message** check box to enable the message box, and enter a customized message. To clear the default message box, select the text and drag down four lines, and press **Delete**.
5. Click **Test** to test your response setup.
6. In the **Select Test Host** dialog box, select the Agent you want to use from the drop-down menu.
7. Click **OK** to apply the Agent and close the **Select Test Host** dialog box.
8. Click **Apply** to update the template.
9. Click **OK** to close the **Template Properties** dialog box.

Setting up SNMP responses

If you select SNMP as a response, Event Monitor uses the SNMP Management Agent to display all the Agents. For example, if a disk fails, Event Monitor sends a trap to the network, which displays the failed disk in a different color.

To set up SNMP responses, be sure the **Single Notification for Multiple Events** check box is not selected on the **General** Response tab.

If you want to configure an IP4700 to send SNMP traps to an NT server, you must install the NT SNMP service on the NT server.

To set up SNMP responses:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template for which you want to set up SNMP responses.
3. Select **Template Properties**.
4. Click the **SNMP** tab.
5. Enter the appropriate information in the **SNMP Management Host** and **Community** fields.
6. Click **Test** to test the Agent.
7. In the **Select Test Host** dialog box, select the Agent you want to use for testing from the drop-down menu.

8. Click **OK** to close the **Select Agent** dialog box.
9. Click **Apply** to update the template.
10. Click **OK** to add the Agent and close the **Template Properties** dialog box.

Receiving SNMP traps

Beginning with Navisphere Agent version 06.19.x.x or higher, SNMP traps sent to your network management station indicate the severity level (see "Understanding SP event codes and event severity" on page 12-34) of the event, using a differentiating severity code in the trap. See the following table for a list of the severity codes and the corresponding severity text.

Severity code	Severity text
3	Information
4	Warning
5	Error
6	Critical (Fault)

If you use Navisphere Integrator with your network management station, you must update Integrator to version 06.19.x.x or higher. The updated version of Integrator registers the severity codes and displays the associated severity text for the event. The severity text displayed by your network management station may differ from the severity levels listed in the table above.

The updated version of Navisphere Integrator continues to support the previous format of the SNMP trap, to maintain compatibility.

Creating templates with customized responses

If the existing responses do not meet your needs, you can create and add customized responses. Event Monitor adds the customized response tab next to the last response tab.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click **Templates** and select **Create New Templates**.

The **Template Properties** dialog box opens.

3. On the **General** tab, select the action that you want the system to use to notify you of events. We recommend **Single Notification for Multiple Events**

4. Click **Add Response**.

The **Add Response** dialog box opens.

5. Type the new response in the **Add Response** field.
6. Click **OK** to add the new response.
Event Monitor adds the new tab to the **Template** dialog box.
7. Enter the appropriate information in the **Program to Execute** and **Program Parameters** fields.
8. Optionally, clear the **Use Default Message** check box to customize the message text.
9. Click **OK**.

Creating new templates based on existing templates

If you have already created templates, Event Monitor lets you create new templates based on existing templates.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click **Templates** and select **Create New Template Based on...**
Or, in the **Tools** menu, select **Event Monitor > Create a New Template Based On...**
3. Select the template that you want to use as a base to create a new template.
4. From the drop-down menu, select the template that you want to use.
5. Click **OK**.

The **Template Properties** dialog box opens, and the new template name appears in the **Template Name** field of the **General** Tab. From here you can change responses for the new template, if desired.

The length of a template filename is limited to 38 characters. In addition, the name must not start with a number, and it must contain characters in the following range: [a...z], [A...Z], [0...9]. If you import an existing template with a filename of illegal length or containing non-supported characters, the file is not replicated across the domain.

Testing responses

Each response tab includes a **Test** button. The test ensures that the response application launches. EMC recommends that you perform a successful test of the response before you actually apply the response template to a host. The test does not detect any typographical errors in the response setup. Be sure that the server you use for the test is the one configured to generate the responses.

To test responses:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template that you want to test and select **Template Properties**.
3. In the **Template Properties** dialog box, click the response tab that you want to test.
4. Click **Test**.

The **Select Test Host** dialog box opens.

5. From the drop-down menu, select the Agent that you want to use to test your response.
6. Click **OK**.

If the test succeeds, the application should be running on the host that runs the specified Agent.

Testing configurations

After you set-up your Event Monitor configurations, you can generate test events, to validate your configurations. Testing your configurations ensures they will work as you expect.

In a centralized monitoring configuration, you can choose to send test events to all monitored Agents within the domain, or you can send a single test event to a monitored Agent you select.

Using insert test event

When you want to generate test events, you can access the insert test event feature by selecting an Agent or storage system icon. If you select a storage system icon, test events are sent to both Agents that reside on the storage system.

In a centralized monitoring environment, if you select the central monitoring Agent, test events are sent to all monitored Agents within the domain. If you select a monitored Agent, a test event is sent only to the selected monitored Agent.

1. In the **Monitors** tab of the **Enterprise Storage** window, right-click the Agent or storage system icon.
2. Click **Insert Test Event**.

The **Insert Test Event** confirmation box opens. The confirmation box lists the hosts to which the test events will be sent.

3. To generate the test events, click **Yes**.

A test event, with a dummy code of 0x2004 and a description of "Test Event," is sent to the selected hosts.

Viewing test events

After you generate test events, the Agent stores the events in the event log. You can view the test events, using the Event Viewer.

For more on viewing events, see "Viewing system event logs" on page 12-36.

Viewing responses to test events

You can view details of the response to test events, in the response log.

For more on viewing responses, see "Viewing the response log" on page 12-37.

Deleting customized responses

When you no longer need a customized response, you can delete it.

You cannot delete the default responses.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template with the customized response and select **Template Properties**.
3. In the **Template Properties** dialog box, click the response tab that you want to delete.
4. Click **Delete**, and click **Yes** to confirm removing the response.

Importing and exporting templates

If you have several domains, you can let them share templates by importing or exporting the templates from a directory path. You can do this with systems that have central monitoring or distributed monitoring.

Importing templates

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click **Templates** and select **Import Template**.

or

In the **Tools** menu, select **Event Monitor** and then select **Import Template**.

3. In the **Import Templates** dialog box, navigate to the folder from which you want to import the templates and click **Select**.

Event Monitor imports the templates to the **Templates** tree in the **Monitors** tab, and updates the Agents with the imported templates.

Exporting templates

1. In the **Enterprise Storage** dialog box, click the Monitors tab.
2. Right-click **Templates** and Select **Export All Templates**.

or

In the **Tools** menu, select **Event Monitor** and then select **Export All Templates**.

3. Click **Yes** to confirm.
4. In the **Export All Templates** dialog box, navigate to the folder to which you want to export all the templates, and click **Select**.

Event Monitor exports all the templates to the selected folder.

Viewing and modifying template properties

Using the right-click menu options in the **Monitors** tab, you can view and modify the properties of a specific template.

When you upgrade a pre-6.2 Agent to a version 6.2 Agent, and then you start the new Agent, the system automatically converts any pre-6.2 templates to the new 6.2 template format. You can modify the converted templates to include the new thresholding properties. Refer to "Adding new threshold properties to converted templates".

Viewing template properties

To view template properties:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template whose properties you want to view, and select **Properties**.

Modifying template properties

To modify template properties:

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template whose properties you want to modify, and select **Properties**.
3. You can modify any of the following properties:
 - In the **General** tab, you can change the events or the severity level of the events you are monitoring.
 - In the **E-Mail**, **Paging** and **SNMP** tabs, you can change the properties of these response methods.

Adding new threshold properties to converted templates

When you upgrade a pre-6.2 Agent to a version 6.2 or higher Agent, the Agent starts automatically, and any pre-6.2 templates are automatically converted to the new 6.2 template format. You can then modify the converted templates to include the new threshold properties.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template whose properties you want to modify, and select **Properties**.
3. In the **General** tab, under **Events**, select **General**, and then click **Advanced** to open the **General Event Threshold Advanced Selection** dialog box.
4. Specify the values for threshold count, threshold interval and response delay for event notification.

Deleting templates

If you no longer need a template, you can delete it. When you delete the template, Event Monitor removes the template from all the storage systems to which it was applied, thereby no longer monitoring the storage systems.

Before you delete your templates, you can save them by exporting them to a specific folder. If you need the templates again, you can easily import them.

Default templates cannot be deleted.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the template that you want to delete, then select **Delete**.

Applying templates to storage systems

If you use the Event Monitor Wizard to configure Event Monitor, you do not need to perform this task.

When you have finished creating templates and selecting the storage systems to be monitored, you can start the monitoring service by applying templates to the storage systems you want to monitor.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the storage system that you want to monitor and select **Monitor Using Template**.

The **Templates** dialog box opens.

3. In the **Select Template** dialog box, select the templates that you want to use to monitor the storage system. Repeat steps 2 and 3 if you want to select multiple Agent.
4. When you are finished, click **OK** to select the templates.

You may experience some delay when the template is downloading to the Agent. This is normal. The default time delay is set to 4 minutes per Agent.

Using Event Monitor

Use the event viewer to view the events of a monitoring host and the events of all hosts it is monitoring. These events include hard errors, startups, and shutdowns involving disks, fans, SPs, LCCs, power supplies, SPSs, and the BBU. The messages are ordered by the time the event occurred with the most recent messages at the beginning of the log. The event viewer has a filtering mechanism built in, so you can view all events or, if you want to view a subset of the events, you can. You can also save any recorded events to the events log file, print a list of the events in the log file, and permanently clear all the events from the log file (see page 12-38).

Understanding SP event codes and event severity

This section explains SP event codes and how they indicate event (error) severity.

There are two different categories of event codes returned by storage system SPs:

- ◆ Basic codes — codes that can occur with any type of storage system. These codes include three or four digits (hexadecimal).
- ◆ Codes from FC4700-Series and CX-Series storage systems only. These codes include eight digits (hexadecimal). The last four digits indicate the specific software module where the condition that provoked the message occurred.

For both categories of events, there are four severity levels. You can select one or more of these levels as phone-home events (or events for which you should be paged or sent an e-mail message), using Navisphere Event Monitor. The severity levels are as follows:

- ◆ Information codes — Require no action by you. Information codes are valuable to engineering in helping to establish history.

Generally, the standard Information codes are three-digit hexadecimal numbers that begin with 6 or 7; for example, 60B and 734. The Event Monitor Information codes are four-digit hex numbers that begin with 2, 4, 3 or 6; for example, 2001 and 6001. Of the four-digit codes that start with 4, the Information codes range from 4*n*00 to 4*n*3F, where the *n* digit identifies the software where the error occurred.

For the CX-Series or FC4700-Series Information codes, the fifth digit is 0. For example, 71050004.

- ◆ Warning codes —Are normal and require no action by you unless they occur frequently.

Generally, the standard Informational Warning codes are three-digit hexadecimal numbers that begin with 8; for example, 805. The Event Monitor Warning codes are four-digit hex numbers that begin with 2, 4, 3 or 6; for example, 2041 and 6041. Of the four-digit codes that start with 4, the Warning codes range from 4 n 40 to 4 n 7F, where the n digit identifies the software where the error occurred.

For the CX-Series or FC4700-Series Warning codes, the fifth digit is 4. For example, 71054003.

- ◆ Error codes — Typically require action by you or a support engineer.

Generally, the standard Error codes are three-digit hexadecimal numbers that begin with 9; for example, 90C. The Event Monitor Error codes are four-digit hex numbers that begin with 2; for example, 2101.

For the CX-Series or FC4700-Series Warning codes, the fifth digit is 8.; for example, 71058003.


- ◆ Critical Error codes —Typically require action by you or a support engineer.

The standard Critical Error codes are three-digit hexadecimal numbers that begin with the hex number a; for example, a02.

For the CX-Series or FC-4700-Series Critical codes, the fifth digit is C. For example, 7215C001.

CX-Series and FC4700-Series code format

Each CX-Series or FC4700-Series message code is an eight-digit hexadecimal number whose meaning is as follows.



For example, in the code 7215C001:

- ◆ 7215 identifies the source module where the event occurred.
- ◆ C specifies the severity (here, Critical Error).
- ◆ 001 identifies the event ID in the module.

Viewing system event logs

Depending on I/O activity, event logs can take time to load and display.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the monitoring host for which you want to view system logs, and select **View Events**.
The **Events** dialog box opens.
3. If you want to view specific event records in the log, click **Filter**.
The **Event Filter** dialog box opens. If you have an event log with several events, you can reduce the log size by filtering the events. Filtering events lets you view only the event types you specify.
4. Enter the appropriate information in the **View**, **View by** and **Types** fields.

5. Click **Default Filter Values** to display all events of all types. When this is checked, the **View**, **Types**, and **View By** fields appear dimmed and are unavailable.
6. When you are finished, click **OK**.

The time stamp in the Event Viewer varies depending on whether the event came from centralized monitoring, distributed monitoring, or the storage system.

Viewing details about a specific event

Event Monitor lets you view detailed information about specific events.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the monitoring host for which you want to view system logs, and select **View Events**.
3. In the **Events** dialog box, double-click the event for which you want to view details.

The **Event Detail** dialog box opens.

4. When you have finished viewing the event detail, you can click **Next**, **Previous**, or **OK**.

If you click **Next**, the details about the next event in the list will be displayed.

If you click **Previous**, the details about the preceding event in the list will be displayed.

5. Click **OK** to close the dialog box.

Viewing the response log

The response log displays a list of all responses sent by this monitored Agent. To see a list of these responses, do the following:

In the **Monitors** tab of the **Enterprise Storage** window, right-click the icon for the monitoring Agent, and then click **View Response Log**.

Viewing the message file

The message file lists any messages that were sent with an event response.

To view the contents of the message file:

1. In the **Monitors** tab of the **Enterprise Storage** window, right-click the icon for the monitoring Agent, and then click **View Message File**.
2. Select a message or messages from the list and then click **OK**.

Manager displays the text of the selected message.

If you selected more than one message, click **OK** in the current message box to display the next message.

Saving events to the log file

The Event Monitor lets you easily save events to the events log file. The events selected by the current filtering are saved.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the monitoring host for which you want to save system logs and select **View Events**.
3. In the **Events** dialog box, click **Save**.
4. In the **Save** dialog box, enter the filename, and click **Save**.

Printing events

Event Monitor lets you easily print events in the log file. However, when you click the **Print** button in the **Events** window, Event Monitor prints all the events in the log file, which can be a large number. We suggest that you print the log file using another application, such as Microsoft Excel or Microsoft Word.

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the monitoring host for which you want to print events and select **View Events**.
3. In the **Events** window, click **Save**.
4. Enter a filename and save as a text (.txt) file.
5. Open the file in another application, such as Microsoft Excel. Highlight only the text that you want to print, and copy. Paste the events on a fresh page, then print.

Clearing events in the events window

Event Monitor lets you easily clear the events in the **Events** window.



CAUTION

Clearing events permanently deletes all the events in the log file.

To clear events

1. In the **Enterprise Storage** dialog box, click the **Monitors** tab.
2. Right-click the monitoring host for which you want to clear events and select **View Events**.
3. In the **Events** dialog box, click **Clear**, and click **Yes** to clear all the events.

What next?

- ◆ Learn how to monitor storage-system operation in Chapter 13, "Monitoring Storage-System Operation"; learn how to reconfigure LUNs, RAID Groups, and Storage Groups in Chapter 8, "Creating LUNs and RAID Groups"; and learn how to reconfigure storage systems in Chapter 15, "Reconfiguring Storage Systems".
- ◆ If your configuration is composed of iSCSI storage systems, refer to Chapter 16, "Managing iSCSI (Internet SCSI) Storage Systems" for information on preparing attached servers and iSCSI storage systems for data I/O.
- ◆ If you will be managing legacy storage systems, refer to Chapter 17, "Managing Legacy Storage Systems and NAS Systems" for information on configuring your domain.

You can monitor the operation of managed storage systems from the Main window. You can also monitor their operation by checking the event messages that the Agent receives from an SP.

For security reasons, we recommend that you place all storage systems in a domain before you configure them.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes the following:

- ◆ Monitoring storage-system operation..... 13-2
- ◆ Displaying storage-system component and server status 13-15
- ◆ Using the SP event viewer for monitoring SP events 13-17
- ◆ Generating and transferring SP log files..... 13-20
- ◆ Monitoring storage-system connectivity status..... 13-21
- ◆ Updating host connectivity status 13-23
- ◆ Monitoring the status of trespassed LUNs..... 13-24
- ◆ Monitoring domain status 13-24
- ◆ Displaying NAS device status..... 13-25

Monitoring storage-system operation

If you are managing many storage systems, it is more convenient to look at the Application icon rather than the storage-system icons. The Application icon is in the **Enterprise Storage** dialog box, on the left side of the title bar. This icon provides overall status of all storage systems managed by the current Manager session.

Each icon in a tree consists of an image representing the component and a description of the component. If the state of the component is not normal, the user will observe an overlay indicating the current state. For example, the letter **T** designates a transitioning object while the letter **F** indicates a faulted object.

Refer to Table 13-1 for descriptions of the icon states.

Table 13-1 Application icon states

Icon	Description
	The component and all of its components are working normally.
	Either the component or one or more of its components has failed.
	Either the component or one or more of its components is in a degraded mode.
	Either the component or one or more of its components is in a transitioning state.

Table 13-1 Application icon states (*continued*)

Icon	Description
	Storage system is unmanaged
	Storage system is inaccessible.
	Storage system is unsupported.

Manual polling

At set intervals, Navisphere Manager polls the storage system and server, and if necessary, updates the status information for the affected storage systems and its components in the tree views in the **Enterprise Storage** window.

Storage system polling

At anytime, you can initiate a manual poll for a storage system by right-clicking the icon for the storage system you want to poll and selecting **Update Now**. Manager refreshes the storage system information displayed in the **Enterprise Storage** trees.

Host polling

At anytime, you can initiate a manual poll for a server by right-clicking the icon for the server you want to poll and selecting **Update Now**. Manager refreshes the server status information.

Storage-system faults

For information about all storage systems with faults, start with step 1 in the following procedure. For information about a specific storage system with faults, start with step 2 in the following procedure.

To monitor the storage-system operation, you can be logged in as global or local administrator, global or local manager, global or local monitor.

To display storage-system faults

1. Access the **Fault Status Report** dialog box for all storage systems by clicking **Faults** from the **Tools** menu on the Main window toolbar.

You can also display the **Fault Status Report** dialog box for a specific storage system by right-clicking the storage-system icon, and then clicking **Faults**.

2. If you want more information about the faulted components in a storage system:
 - a. From the **Storage** tab in the **Enterprise Storage** dialog box, double-click the icon for the faulted storage system to expand the **Storage** tree.
 - b. For each faulted component icon that has a menu associated with it, examine its properties as follows:
 - Right-click the faulted icon and click **Properties**.
 - If the dialog box has different tabs, click a tab to view the additional properties.
 - c. For each faulted icon that does not have a menu associated with it, do the following:
 - Double-click the icon to display the icons for its components.

- d. For more information about a FRU represented by a faulted or icon, go to the selection listed below for the FRU.

For a faulted	Go to the section
Disk icon	<i>Faulted disk icon</i> on page 13-5.
SP icon	"Faulted SP icon" on page 13-8.
LCC icon	"Faulted LCC icon" on page 13-9.
Fan A or Fan B icon	"Faulted fan A or fan B icon" on page 13-10.
Power supply or VSC	"Faulted LCC icon" on page 13-9.
SPS icon	"Faulted LCC icon" on page 13-9.
BBU icon	"Faulted BBU icon" on page 13-14.

Faulted disk icon

A faulted disk icon indicates that the disk it represents is in one of the states listed in Table 13-2.

Table 13-2 Disk failure states

State	Meaning
Removed	Removed from the enclosure; applies only to a disk that is part of a LUN.
Off	Failed and powered off by the SP.

You can determine the state of a disk from the **General** tab of its **Disk Properties** dialog box (page 13-16).



CAUTION

Removing the wrong disk can introduce an additional fault that shuts down the LUN containing the failed disk. Before removing a disk, be sure to verify that the suspected disk has actually failed by looking at its orange check or fault light or the SP event log for the SP that owns the LUN containing the disk. In addition to viewing the log for messages about the disk, also look for any other messages that indicate a related failure, such as a failure of a SCSI bus or a general shutdown of an enclosure. Such a message could mean the disk itself has not failed. A message about the disk will contain its module ID.

The icon for a working hot spare in a RAID Group may be orange instead of blue if you replace the failed disk that the hot spare is replacing while the hot spare is transitioning into a group. When this happens, the icon for a working SP is orange instead of green. The Fault Status Report dialog box reports that the storage system is normal instead of transitioning, the state property for the hot spare is faulted instead of transitioning, and the state of the SP is normal (the correct state).

After you confirm the failure of a disk, the system operator or service person should replace it, as described in the storage-system installation and service manual.

You must replace a failed disk with one of the same capacity and format.

The remainder of this section discusses a failed disk in a RAID 0 or disk LUN or a RAID 5, 3, 1, or 1/0 LUN and a failed vault disk when storage-system write caching is enabled.

Failed disk in a RAID 0 or disk LUN

If a disk in a RAID 0 or disk LUN fails, applications cannot access the LUN.

Before you replace the failed disk, unbind the LUN.

After you replace the failed disk:

1. Rebind the LUN.
2. Create partitions or file systems on the LUN before restoring data from backup files.

Failed disk in a RAID 5, 3, 1, or 1/0 LUN

If a disk in a RAID 5, 3, 1, or 1/0 LUN fails, applications can continue to access the LUN. If the storage system has a hot spare on standby when a disk fails, the SP automatically rebuilds the failed disk on the hot spare.

You should replace the failed disk while the storage system is powered up so that applications can continue to access the LUN. When you replace the disk, the SP:

- ◆ Formats and checks the new disk.
- ◆ Rebuilds the data on the new disk as described below.

While rebuilding occurs, applications have uninterrupted access to information on the LUN.

Rebuilding a RAID 5, 3, 1, or 1/0 LUN

You can monitor the rebuilding of a new disk from the **General** tab of its **Disk Properties** dialog box (page 13-16).

A new disk module's state changes as follows:

1. **Powering up** - The disk is powering up.
2. **Rebuilding** - The storage system is reconstructing the data on the new disk from the information on the other disks in the LUN.

If the disk is the replacement for a hot spare that is being integrated into a redundant LUN, the state is Equalizing instead of Rebuilding. In this situation, the storage system is simply copying the data from the hot spare onto the new disk.

3. **Enabled** - The disk is bound and assigned to the SP being used as the communication channel to the enclosure.

A hot spare's state changes as follows:

1. **Rebuilding** - The SP is rebuilding the data on the hot spare.
2. **Enabled** - The hot spare is fully integrated into the LUN, or the failed disk has been replaced with a new disk and the SP is copying the data from the hot spare onto the new disk.
3. **Ready** - The copy is complete. The LUN consists of the disks in the original slots and the hot spare is on standby.

Rebuilding occurs at the same time as user I/O. The rebuild priority for the LUN determines the duration of the rebuild process and the amount of SP resources dedicated to rebuilding.

A High or ASAP (as soon as possible) rebuild priority consumes many resources and may significantly degrade performance. A Low rebuild priority consumes fewer resources with less effect on performance. You can determine the rebuild priority for a LUN from the **General** tab of its **LUN Properties** dialog box (page 13-16).

Failed vault disk with storage-system write caching enabled

If you are using write caching, the storage system uses the disks listed in Table 13-3 for its cache vault. If one of these disks fails, the storage system dumps its write cache image to the remaining disks in the vault; then it writes all dirty (modified) pages to disk and disables write caching.

Storage-system write caching remains disabled until a replacement disk is inserted and the storage system rebuilds the LUN with the replacement disk in it. You can determine whether storage-system write caching is enabled or disabled from the **Cache** tab of its **Properties** dialog box (page 13-15).

Table 13-3 Cache vault disks

Storage-system type	Cache vault disks
CX-Series	0-0 through 0-14
FC4400/4500, FC47000 FC5600/5700	0-0 through 0-8
FC5200/5300	0-0 through 0-4
C1900, C2x00, C3x00	A0, B0, C0, D0, E0
C1000	A0 through A4

Faulted SP icon

A faulted SP icon indicates that the SP it represents has failed. When an SP fails, one or more LUNs may become inaccessible and the storage system's performance may decrease. In addition, the SP's check or service light turns on, along with the check or service light on the front of the storage system.

If the storage system has a second SP and failover software is running on the server, the LUNs that the failed SP owned may be accessible through the working SP. If the server is not running failover software and a number of LUNs are inaccessible, you may want to transfer control of the LUNs to the working SP (Chapter 14).



CAUTION

The icon for a working SP may appear faulted instead of normal when you replace a failed disk in a RAID Group with a working hot spare while it is transitioning into the group to replace the failed disk. When this happens, the icon for the hot spare is faulted instead of transitioning. The Fault Status Report dialog box reports that the storage system is normal instead of transitioning, the state property for the SP is normal (the correct state), and the state property for the hot spare is faulted instead of transitioning.

The system operator or service person can replace the SP under power; however, doing so interrupts application access to LUNs owned by the SP unless the server is running LUN transfer software such as PowerPath. If an SP in a shared storage system needs replacing, refer to the service provider's installation and service manual for switch configurations. If an SP in an unshared storage system needs replacing, refer to the storage-system installation and service manual for replacement instructions.

Faulted LCC icon

A faulted link control card (LCC) icon indicates that the LCC it represents has failed. In addition, the LCC's fault light turns on, along with the service light on the front of the storage system.

When an LCC fails, the SP it is connected to loses access to its LUNs, and the storage system's performance may decrease. If the storage system has a second LCC and the server is running failover software, the LUNs that were owned by the SP connected to the failed LCC may be accessible through the SP connected to the working LCC. If the server is not running failover software, you may want to transfer control of the inaccessible LUNs to the SP that is connected to the working LCC (Chapter 14).

The system operator or service person can replace the LCC under power, without interrupting applications to accessible LUNs. The storage-system installation and service manual describes how to replace an LCC.

If the LCC is removed from Bus X, Enclosure Y, the Agent can no longer gather information along this path, and all devices along this path will show as faulted.

Faulted fan A or fan B icon

For any FC-Series storage system, a faulted Fan A icon in enclosure n (DAEs) indicates that the drive fan pack has one or more faults. A faulted Fan A or B icon in enclosure 0 (DPE) indicates that the SP fan pack has one or more faults.

For CX-Series storage systems, a faulted Fan A or B icon in enclosure n (DAEs) indicates that the drive fan pack has one or more faults. A faulted Fan A, B, or C icon in the SPE enclosure indicates that the SP fan pack has one or more faults.

Only SPEs for CX-Series and DPEs for FC-Series have SP fan packs.

For any C-series storage system, a faulted Fan A icon and a green and grey Fan B icon indicate that its fan module has one fault. A faulted Fan A icon and a faulted Fan B icon indicate that its fan module has two or more faults.

Drive fan pack

If one fan fails in a drive fan pack, the other fans speed up to compensate so that the storage system can continue operating. If a second fan fails and the temperature rises, the storage system shuts down after about two minutes.

If you see a faulted Fan A icon in an FC-series storage system, the system operator or a service person should replace the entire drive fan pack as soon as possible. The storage-system installation and service manual describes how to replace the fan pack.

Do not remove a faulted drive fan pack until a replacement unit is available. You can replace the drive fan pack while the DPE or DAE is powered up.

If the drive fan pack in a DPE is removed for more than two minutes, the SPs and the disks power down. The SPs and disks power up when you reinstall a drive fan pack.

If the drive fan pack in a DAE is removed for more than two minutes, the Fibre Channel interconnect system continues to operate, but the disks power down. The disks power up when you reinstall a drive fan pack.

SP fan pack If one fan fails in an SP fan pack, the other fans speed up to compensate so that the storage system can continue operating. If a second fan fails and the temperature rises, the storage system shuts down after about two minutes.

If you see a faulted Fan B icon, the system operator or a service person should replace the entire fan pack or module as soon as possible. The storage-system installation and service manual describes how to replace the fan pack or module.

Do not remove a faulted SP fan pack until a replacement unit is available. You can replace the fan pack when the DPE is powered up. If the fan pack is removed for more than two minutes, the SPs and the disks power down. They power up when you reinstall an SP fan pack.

Fan module Each C-series storage system has one fan module. Table 13-4 shows the number of fans per number of slots in the enclosure.

Table 13-4 Number of fans in a C-series storage system

Enclosure size	Number of fans
30-slot	9
20-slot	6
10-slot	3

If any fan fails, the fault light on the back of the fan module turns on. The storage system can run after one fan fails; however, if another fan failure occurs, the storage system shuts down after two minutes.

If you see a faulted Fan A icon in a C-series storage system, the system operator or a service person should replace the entire fan module as soon as possible. The storage-system installation and service manual describes how to replace the fan module.

Swinging the fan module away from the enclosure or removing it for more than two minutes may cause the storage system to overheat. To prevent damage to the disks from overheating, the storage system shuts down if you unlatch or remove the fan module for more than two minutes. You should not leave a fan module unlatched or removed for more than the absolute minimum amount of time that you need to replace it.

Faulted power supply or VSC icon

A faulted power supply or voltage semi-regulated converter (VSC) icon indicates that the power supply it represents has failed.

CX-Series and FC-Series Storage System

Each enclosure has one or two power supplies: A and optionally B. An enclosure with two power supplies can recover from the failure of one power supply and provide uninterrupted service while the defective power supply is replaced. If a second power supply fails or is removed, the enclosure shuts down immediately.

C3x00 Series Storage System

The system has three power supplies (VSCs): A, B, and C. It can recover from the failure of one power supply and provide uninterrupted service while the defective power supply is replaced. If a second power supply fails or is removed, the enclosure shuts down immediately.

C2x00 Series Storage System

The system has two or three power supplies: A, B, and optionally C. If it has three power supplies, it can recover from the failure of one power supply and provide uninterrupted service while the defective power supply is replaced. If a second power supply fails or is removed, the enclosure shuts down immediately.

C1900 or C1000 Series Storage System

The system has one or two power supplies: A and optionally B. If it has two power supplies, it can recover from the failure of one power supply and provide uninterrupted service while the defective power supply is replaced. If a second power supply fails or is removed, the enclosure shuts down immediately.

Storage System Without an SPS or BBU

Failure of the ac distribution system (line cord, utility power, and so on) also immediately shuts down the entire enclosure.

When a power supply fails, the system operator or service person should replace the power supply as soon as possible. The storage-system installation and service manual describes how to replace a power supply.

When a C-series enclosure or the DPE in an FC-series storage system shuts down, the operating system loses contact with the LUNs. When the enclosure powers up, you may need to reboot the server to let the operating system access LUNs, and you must restart the Agent on the server connected to the storage system.

Faulted SPS icon

A faulted standby power supply (SPS) icon indicates that the SPS it represents has an internal fault. When the SPS develops an internal fault, it may still be able to run on line, but the SPs disable write caching. The storage system can use the write cache only when a fully charged, working SPS is present.

However, if the storage system has a second SPS, write caching can continue when one SPS has an internal fault or is not fully charged. The status lights on the SPS indicate when it has an internal fault, when it is recharging, and when it needs replacing because its battery cannot be recharged.

Each week, the SP runs a battery self-test to ensure that the monitoring circuitry is working in each SPS. While the test runs, storage-system write caching is disabled, but communication with the server continues. I/O performance may decrease during the test. When the test is finished, storage-system write caching is re-enabled automatically. The factory default setting for the battery test is set to start at 1:00 a.m. on Sunday, which you can change (Chapter 7).

When the SPS Fault light or the SPS Replace Battery light is lit, the system operator or service person should replace the SPS as soon as possible. The SPS installation and service manual describes how to replace an SPS.

If the storage system has two SPSs, you can replace one of them while the DPE or SPE is powered up, but we recommend that you disable storage-system write caching before removing the SPS (Chapter 7).

Faulted BBU icon

A faulted battery backup unit (BBU) icon indicates that the BBU in a C-series storage system is in one of the states listed below.

Table 13-5 BBU failure states

BBU state	Meaning
Down	Failed or removed <i>after</i> the Agent started running
Not Present	Failed or removed <i>before</i> the Agent started running

You can determine the state of a BBU from its **Properties** dialog box (page 13-16).

When a BBU fails:

- ◆ Storage-system write caching is disabled and storage-system performance may decrease.

You can determine whether storage-system write caching is disabled from the **Cache** tab of the storage-system **Properties** dialog box (page 13-15).

Storage-system write caching remains disabled until the BBU is replaced.

- ◆ The BBU service light turns on, indicating that the BBU is either charging or not working.

After a power outage, a BBU takes 15 minutes or less to recharge. From total depletion, recharging takes an hour or less.

Each week, the SP runs a self-test to ensure that the BBU's monitoring circuitry is working. While the test runs, storage-system caching is disabled, but communication with the server continues. I/O performance may decrease during the test. When the test is finished, storage-system caching is re-enabled automatically. The factory default start time for the BBU test is 1:00 a.m. on Sunday, which you can change (Chapter 7).

A system operator or service person can replace a failed BBU under power without interrupting applications. The storage-system installation and service manual describes how to replace the BBU.

Displaying storage-system component and server status

Most hardware components, each storage-system server, each RAID Group, each LUN, and each Storage Group represented by an icon on the **Storage** or **Hosts** tree have a **Properties** dialog box associated with them. The **Properties** dialog box provides a variety of information about the component or simply displays the current state of the component.

To monitor the storage-system operation, you can be logged in as global or local administrator, global or local manager, global or local monitor.

For a description of each property in a dialog box, click **Help**.

To display storage-system properties

From the **Storage** tab of the **Enterprise Storage** dialog box, right-click the icon for the storage system, and click **Properties**.

The **Storage System Properties** dialog box opens with the **General** tab displayed. You can also display properties for the storage system's cache, memory, storage access and software.

To display storage-system server properties - FC4400/4500 only

From the **Host** tab of the **Enterprise Storage** dialog box, navigate to the icon for server for whose properties you want to display, and then select **Properties**.

The **Host Properties** dialog box opens with the **General** tab displayed.

For a description of each property, click **Help** in the dialog box.

If the server is connected to a shared storage system and you want to view the server's storage properties, select the **Storage** tab.

To display storage group properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group whose properties you want to display and select **Properties**.

To display LUN properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the LUN whose properties you want to display, and then select **Properties**.

To display SP properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the SP whose properties you want to display, and then select **Properties**.

To display RAID group properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group whose properties you want to display, and then select **Properties**.

To display disk properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the disk whose properties you want to display, and then select **Properties**.

To display SPS or BBU properties

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the SPS or BBU whose properties you want to display, and then select **Properties**.

To display LCC state

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the LCC whose properties you want to display, and then select **Properties**.

To display power supply state

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the power supply whose state you want to display, and then select **Properties**.

Using the SP event viewer for monitoring SP events

Each SP maintains a log of event messages. These events include hard errors, startups, and shutdowns involving disks, fans, SPs, LCCs, power supplies, SPSs, and the BBU. The messages are ordered by the time the event occurred with the most recent messages at the beginning of the log. The event viewer has a filtering mechanism built in, so you can view all events or, if you want to view a subset of the events, you can.

For an overview of event code format see "Understanding SP event codes and event severity" on page 12-34.

Displaying the event logs for an FC4700 SP

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Double-click the icon for the storage system with the SP whose event log you want to display.
3. Right-click the icon for the SP whose event log you want to display, and click **Event Log**.

You can also open the event logs from the **Event Monitor Configuration** window. To do so, right-click the icon for the monitoring host (SP Agent) for which you want to view the storage-system logs, and click **View Events**. This displays the events of the selected monitoring host and the events of all hosts it is monitoring.

The fields in the **Events** window have the following meanings:

Date - Day the event occurred.

Time - Time the event occurred.

Event Code - The hexadecimal code for the type of event that occurred.

Description - An abbreviated description of the message. See the *Storage-System and Navisphere Event Messages Reference* for a more detailed description of codes.

Storage System- The name of the storage system that the event is about.

SP - The name of the SP that the event is about.

Host - Name of the host that the Agent is running on.

Filtering events

If you have an event log with multiple events, you can reduce the log size by filtering the events. Filtering events lets you view only the event types you specify.

1. In the **Events** window, click **Filter**.
2. Under **View**, select one of the options from the **As of** list.
3. Under **Types**, select the types of events that you want to view.
4. Under **View By**, select the event code by which you want to view the events, or in **Description**, enter a specific event description.

For example, if you want to track activities to users, select the **0x4600-0x46ff audit log** event code, and click **OK**.
5. Click **OK**.

Viewing event details

1. In the **Events** window, double-click the event for which you want to view details.
2. When you have finished viewing the event detail, you can do any of the following:
 - Click **Next** to view the next event detail.
 - Click **Previous** to view the previous event detail.
 - Click **Close** to close the **Event Detail** dialog box.
3. For more information about the properties in the dialog box, click **Help**.

Saving events to a log file

1. In the **Events** window, click **Save**.
A **Save As** dialog box opens.
2. In **File name**, enter the name of the file in which you want to save the events displayed in the **Events** window.
3. Click **Save**.

Printing events

You can print all the events displayed in the Event window by clicking the **Print** button in the window. Since the number of displayed events may be every large, we recommend that you save the events to a file, and print the file using another application, such as Microsoft Excel or Microsoft Word, as follows:

1. In the **Events** window, click **Save**.
A **Save as** dialog box opens.
2. In **File name**, enter the name of the file in which you want to save the events displayed in the **Events** window.
3. In **Save as type**, select **Text Files (*.txt)** from the list.
4. Open the file in another application, such as Microsoft Excel or Notepad.
5. Highlight only the text that you want to print, and copy the text to the clipboard.
6. Paste the events on a fresh page in the application.
7. Print your file.

Clearing events in the events window



CAUTION

Clearing events permanently deletes all the events in the log file.

1. In the **Events** window, click **Clear**.
2. Click **Yes** to clear all the events from the **Events** window.

Generating and transferring SP log files

If the storage system is experiencing problems, you can generate log files for the storage system's SPs that will help to diagnose the problems. Once the log files are generated, transfer them from the storage system to a local client so that they are available for review by technical support.

1. In the **Enterprise Storage** dialog box, click the **Storage** tab.
2. Expand the icon for the storage system that owns the SP for which you want to generate log files.
3. Right-click the icon for the SP and click **SP Collect**.

Navisphere generates a zip file that includes the log files. The format of the zip file is

SPx_arrayserialnumber_spsignature_date_time_data.zip

where

x is A or B

arrayserialnumber is the storage system's serial number.

spsignature is the SP's signature.

date is the date the data was collected.

time is the time the data was collected.

An example of a valid entry is

SPA_APM00043705741_1198138_02-09-2005_20-51-43_data.zip

4. Transfer the zip file to a directory that is not on the storage system.
 - a. Right-click the icon for the SP and click **File Transfer Manager**.
 - b. In **Destination Directory**, click **Browse** to select where you want to transfer the zip file.
 - c. In the **Files on SP** list, select the data zip file that you want to transfer and click **Transfer**.

The **Transfer Status** box updates the status of the file transfer operation.

If you click **Cancel** while a file is being transferred, the file transfer will complete, but any other selected files will *not* be transferred.

Monitoring storage-system connectivity status

To monitor and configure the storage system connectivity status, you must be logged in as global or local administrator, or global or local manager.

The storage system **Connectivity Status** dialog box lets you monitor and configure the front-end activity (HBA to SP Port and SP Port to SP Port) for the storage system. The **Connectivity Status** dialog box includes an entry for each HBA or SP Port connected to the selected storage system. The listed HBAs are either currently logged in or permanently registered with the storage system.

For the HBA to be logged in, the following are required:

- ◆ The HBA must be powered up.
- ◆ There must be a connection from the server to the storage system.
- ◆ The FLARE software must recognize the HBA.

For the HBA to be permanently registered, the Agent must recognize the HBA, and when all HBAs belonging to a server are registered, the server can be connected to and disconnected from Storage Groups.

If you remove and then replace an HBA in a server, right-click the server in which you replaced the HBA and click **Refresh Initiators**. This will enable the server to see all LUNs belonging to Storage Groups owned by either SP A, SP B or both. This eliminates the need to disconnect the server from and then reconnect the server to the Storage Group.

For each initiator listed in the dialog box, you can selectively view, add and modify initiator information. This initiator information includes:

- ◆ Initiator Type
- ◆ Communication Path from the initiator to the SP (Array CommPath)
- ◆ Failover Mode
- ◆ The name and IP Address of the server in which the initiator resides
- ◆ Unit Serial Number

For a description of each property in the dialog box, click **Help**.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the storage system for which you want to monitor or configure the connectivity status, and choose **Connectivity Status**.
2. To show or hide any connections from the SP Ports of other storage systems to the SP Ports of the selected storage system, select or clear **Show Storage System Initiators**.
3. To register any of the listed initiators, select the initiators, click **Register**, and then enter the initiator information.
4. To create a new initiator record, click **New**, and then enter the initiator information.
5. To display information about a specific initiator, select the initiator, and click **Info**.
6. When you physically disconnect an HBA from a storage system, the initiator record for that HBA still displays in the **Connectivity Status** dialog box. To remove this record, select it and then click **Deregister**.

If the initiator that you select is logged in when you click **Deregister**, Manager will *not* delete it from the **Initiator Records** list. If it was not logged in, Manager will permanently delete it from the list.

7. To edit the initiator information for a group of initiators, click **Group Edit**, select the initiators you want to edit, and then enter or modify the initiator information. See the online help for more information.

Updating host connectivity status

The **Host Connectivity Status** dialog box, lets you maintain connectivity to your storage system after you replace a server HBA, without having to disconnect the server and then reconnect it to Storage Groups. The reconnect feature updates all Storage Groups that are connected to the server with any new HBA information.

You can use the reconnect feature on only one server at a time.

Requirements

Before you can use the reconnect feature, the following must be true:

- ◆ Each HBA in the server is zoned to a different storage system SP. For example, HBA 1 is zoned to SP A, and HBA 2 is zoned to SP B.
- ◆ Any new or replaced HBAs have been registered with the storage system.
- ◆ The server belongs to a storage-system Storage Group.

To update server connectivity

1. In the **Hosts** tree, right-click the name of the server for which you want to update connectivity status (server with new or replaced HBAs), and click **Connectivity Status**.

The **Connectivity Status** dialog box lists all the storage-system Storage Groups to which this server has access.

2. Click **Reconnect** to start the reconnect operation.

Navisphere Manager will display a confirmation dialog box listing the Storage Groups to which the reconnect operation will be attempted.

3. Click **Yes** to continue the reconnect operation; click **No** to cancel the operation and return to the **Host Connectivity Status** dialog box.

The performance of the reconnect operation depends on the number of Storage Groups that need to be updated with the new HBA information.

Monitoring the status of trespassed LUNs

The **Trespassed LUNs** dialog box lets you display information and status about all trespassed LUNs (LUNs whose current SP ownership is not the default SP) assigned to a server. You can display the following information about trespassed LUNs:

- ◆ Host Name
- ◆ Device ID
- ◆ Trespassed LUN ID
- ◆ Current SP
- ◆ Default SP

Click **Help** in the dialog box to display descriptions of each property.

To open the **Trespassed LUNs** dialog box, do the following:

- ◆ On the **Tools** menu, select **Trespassed LUNs**.

Monitoring domain status

The **Domain Status** dialog box displays status information for *all* systems within the domain that are running Storage Management Server software. It does not display the status of any legacy systems being managed through a portal.

1. In the **Enterprise Storage** dialog box, in the **File** menu, select **Setup Domain**, and then click **Domain Status**.

The **Domain Status** dialog box provides the following information about each system running Storage Management Server software.

For a detailed description of each field in the dialog box, click **Help**.

- System name
- IP addresses for each system
- Status of the connection between the client and the Storage Management Server software
- Communication ports

- Platform type - Describes whether the system is a storage system or a server running Storage Management Server software.
- Storage Management Server software revision
- Status as a portal - Yes or No

Displaying NAS device status

You can determine the status of a NAS device by displaying its properties.

To display NAS device properties

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the NAS device and click **Properties**.
2. To display information about a network interface connection to the NAS device, click one of the **Network Interface** tabs.
3. For a description of each property in the dialog box, click **Help**.

Reconfiguring LUNs, RAID Groups, and Storage Groups

After you create a LUN, RAID Group or Storage Group you may want to reconfigure it by changing its properties.

For security reasons, we recommend that you place all storage systems in a domain before you configure them.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

This chapter describes

- ◆ Reconfiguring LUNs..... 14-2
- ◆ Using the LUN migration feature to copy a LUN..... 14-13
- ◆ Reconfiguring RAID groups..... 14-17
- ◆ Reconfiguring storage groups..... 14-23

Reconfiguring LUNs

To reconfigure the storage system you must be logged in as global or local administrator, or global or local manager.

After you bind a LUN, you can change all of the LUN's properties without unbinding it (and thus losing its data), except for the following:

- ◆ Unique ID
- ◆ Element size
- ◆ RAID type

To change any of these properties, follow the procedures below.



CAUTION

Unbinding a LUN destroys any data on the LUN. Before unbinding a LUN, make a backup copy of any data that you want to retain. Do not unbind the last LUN owned by an SP connected to a NetWare or Solaris server unless it is absolutely necessary. If you do unbind it, you will have to do the following:

NetWare server - Refer to the Release Notice for the NetWare Navisphere Agent for information on how to bind the first LUN.

Solaris server - Edit the Agent configuration file on the servers connected to the SPs that you will use to bind LUNs.

To change a LUN's unique ID or element size

You cannot change the ID or element size of a LUN without unbinding the LUN (and thus losing all its data) and then binding a new LUN with the desired ID or element size.

To change the RAID type of a LUN

Unbind the LUN, then do one of the following:

- ◆ For a non-RAID Group storage system, bind a new LUN with the desired RAID type.
- ◆ For a RAID Group storage system, bind a new LUN on another RAID Group that supports the desired RAID type.

To change the user capacity of a LUN

In most cases, you cannot change the user capacity of a LUN without unbinding the LUN and losing all data on it. The only exception is if the LUN is the only LUN within a RAID Group and if its user capacity equals the user capacity of the RAID Group. For such a LUN, you can increase its user capacity by expanding its RAID Group as described on page 14-18.

In a RAID Group, you can rebind a LUN with an increased user capacity only if the Group has usable space that is at least as great as the user capacity you want to add. To determine the amount of usable space in a RAID Group, refer to page 14-20.

If the RAID Group does not have enough usable space, you may be able to free up enough space by defragmenting the RAID Group as described on page 14-19. If defragmenting does not work, then you must expand the RAID Group as described on page 14-18.

All disks in a LUN in a non-RAID Group storage system must have the same physical capacity to fully use the storage space on the disks.

The physical capacity of a hot spare LUN must be at least as great as the physical capacity of the largest disk module in any LUN or RAID Group on the storage system.

Before changing the user capacity of a LUN, be sure to back up any data you want to retain on the LUN. How you change the user capacity of a LUN depends on whether the LUN is part of a RAID Group and what its user capacity is.

If the RAID Group does not have the additional user capacity you want for a LUN, you must first expand the RAID Group.

- ◆ To change the user capacity of a LUN if the LUN is the only one on a RAID Group and its user capacity is equal to the user capacity of the RAID Group:
 - Expand the RAID Group, and the LUN automatically increases its user capacity to the size of the expanded RAID Group.
- ◆ To change the user capacity of a LUN if the LUN is on a non-RAID Group storage system:
 - Unbind the LUN, and then bind a new LUN with more or less disks, depending on whether you want to increase or decrease the user capacity.
 - Make the LUN available to the operating system.
- ◆ To change the user capacity of a LUN that is on a RAID Group storage system:
 - Unbind the LUN, and then bind a new LUN with either additional or less user capacity.
 - For shared storage systems, add the new LUN to the original Storage Group and then make the LUN available to the operating system.
 - For unshared storage systems, make the new LUN available to the operating system.

Changing the remaining properties of a LUN

Changing the following properties of a LUN does not affect the data stored on the LUN:

- ◆ Read caching

A LUN with read caching enabled uses default values for its prefetching properties. The next section describes how to change these properties.

A LUN with read caching enabled can use read caching only if the read cache for the SP that owns it is enabled. Similarly, a LUN with write caching enabled can use write caching only if the storage-system write cache is enabled. To enable the read cache for an SP or the storage-system write cache, refer to Chapter 7.
- ◆ Write caching
- ◆ Rebuild Priority

- ◆ Verify Priority
- ◆ Auto Assign

To change any of these properties, do the following:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, select the icon for the LUN whose cache properties you want to change, right-click, and select **Properties**.
2. For read caching and write caching, select the **Cache** tab to locate the properties you want to change.
3. For the other properties, select the **General** tab to locate the properties you want to change.

Changing LUN prefetch (read caching) properties

Prefetching is read-ahead caching. This process lets the SP anticipate the data that an application will request so it can read the data into its read cache before the data is needed. The SP monitors I/O requests to each LUN that it owns for sequential reads. If it finds that any I/O requests occur, it prefetches the data for them automatically from the LUN.

You can define a specific type of prefetch operation for any LUN, except a RAID 3 LUN or hot spare, by setting the values of the LUN's prefetch properties. Changing any of these properties does not affect the data on the LUN.

Prefetch type - Determines whether to prefetch data of a variable or constant length or disable prefetching.

Prefetch size or prefetch multiplier - Determines the amount of data prefetched for one host read request.

For constant-length prefetching, the prefetch size is the number of data blocks to prefetch. For variable-length prefetching, the prefetch multiplier is the amount of data to prefetch relative to the amount of data requested. For example, if the prefetch multiplier is 8, the amount of data to prefetch is 8 multiplied by the amount of data requested.

Segment size or segment multiplier - Determines the size of the segments that make up a prefetch operation. The SP reads one segment at a time from the LUN because smaller prefetch requests interfere less with other host requests.

For constant-length prefetching, the segment size is the number of data blocks to prefetch in one read operation from the LUN. For variable-length prefetching, the segment multiplier determines the amount of data to prefetch in one operation relative to the amount of data requested. For example, if the segment multiplier is 4, the segment size is 4 multiplied by the amount of data requested.

Maximum prefetch - Determines the number of data blocks to prefetch for variable-length prefetching.

Retention - Determines whether prefetched data has equal or favored priority over host-requested data when the read cache becomes full.

Idle count - With prefetching enabled, specifies the maximum number of I/Os that can be outstanding to the storage system. When this number is exceeded, the system disables prefetching.

Disable size - Determines when a read request is so large that prefetching data would not be beneficial. If so, the read request is disabled. For example, if the amount of requested data is equal to or greater than the size of the read cache, prefetching is a waste of resources.

Table 14-1 Default prefetch properties values

Property	Default value
Prefetch type	Variable
Prefetch multiplier	4
Segment multiplier	4
Maximum prefetch	512 blocks*
Retention	Favor prefetch
Idle count	40
Disable size	129 blocks
* For CX-Series storage systems, the default Maximum Prefetch value is 4096.	

We recommend that you use the default values, unless you are certain that the applications accessing the LUN will benefit from changing the values.

Table 14-2 Available prefetch properties values - general

General prefetch properties	Valid values
Prefetch type	None, constant, variable
Retention	Favor prefetch, Equal priority
Idle count	0 through 100
Disable size	0 through 65534 blocks

Table 14-3 Available prefetch properties values - constant

Constant prefetch properties	Valid values
Prefetch size	0 through 2048 blocks and equal to or greater than the segment size
Segment size	0 if prefetch size is 0 1 through prefetch size with a maximum of 254

Table 14-4 Available prefetch properties values - variable

Variable prefetch properties	Valid values
Prefetch multiplier	0 through 32
Segment multiplier	0 if prefetch multiplier is 0 1 through prefetch multiplier
Maximum prefetch	0 through 2048 blocks*
* For CX-Series storage systems, the available Maximum Prefetch value is 8192.	

To change prefetch properties

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the LUN whose prefetch properties you want to change, right-click, and then select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. In the **Prefetch** tab, set the prefetch properties by doing one of the following:
 - If you want to use the default values, select the **Use Default Values** check box.
 - If you do not want to use the default values, make sure the **Use Default Values** check box is not selected, and make the changes to the prefetch properties.
3. Click **OK** to save the settings.

Transferring the default ownership of a LUN

In a storage system, the default owner is the SP that assumes ownership of a LUN after storage-system power is turned off and then on again.

In the storage system, you can transfer the default ownership of a LUN from the SP that is the default owner of the LUN (primary route to the LUN) to the other SP (secondary route to the LUN).

Transferring default ownership of a LUN is one way of transferring control of a LUN.

You should transfer the default ownership of a LUN when you want to balance LUNs between two SPs, as you might, for example, if a second SP is installed in a storage system with existing LUNs.

Depending on the type of server connected to the storage system, you may want to transfer default ownership of a LUN if any of the following failure situations occurs:

- ◆ An SP fails and the server to which it is connected is not running failover software, and you want to transfer control of a LUN to the working SP.
- ◆ Each SP is connected to a different host bus adapter and one adapter fails or the connection to one adapter fails, and you want the working adapter to access the LUNs owned by the SP connected to the other adapter.

- ◆ When one server fails in a dual-server configuration without host failover software, and you want the working server to access the failed server's LUNs.

The auto assign property of a LUN and failover software can also transfer control of a LUN from one SP to another. For information on the auto assign property, refer to Chapter 7. If you have failover software on the server, you should use it to handle the failure situations just listed, instead of the procedure in this section.

Transferring default ownership of a LUN from one SP to another can affect how the operating system accesses the LUN. Any change you make in ownership does not take effect until the storage system is powered down and up again.

To transfer default ownership of a LUN

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the LUN whose prefetch properties you want to change, right-click, and then select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. In the **General** tab, under **Default Owner**, click **SP A** to make SP A the default owner of the LUN, or click **SP B** to make SP B its default owner.
3. If the storage system is connected to a Solaris server, unmount all the partitions that are associated with the LUN.
4. Have the system operator or service person power the storage system down and then up again for the change in ownership to take effect.
5. On each server with access to the LUN whose SP ownership you transferred, follow the procedure below for the operating system on the server:

AIX:

- a. Remove the **sp** device for the LUN on the original SP using the **rmdev -l sp# -d** command.
- b. Rescan the bus to make AIX aware of the transferred LUN using the **cfgmgr** command.

NetWare:

- a. Scan all buses to make NetWare aware of the transferred LUN using the **scan all luns** command.
- b. Verify that NetWare sees the transferred LUN using the **list devices** command.

HP-UX:

- a. Make sure that the storage system has finished re-initializing and that the LUN has transferred.
- b. Rescan the bus and inform HP-UX about the transferred LUN using the **ioscan -fnC disk** command.

Solaris:

- a. LUN identifies the SP that is the new default owner and that this device name is in the **/kernel/drv/sd.conf** file.
- b. Shut down Solaris.
- c. Restart Solaris using the **boot -r** command.

Windows:

Reboot Windows.

Unbinding a LUN

Typically, you unbind a LUN only if you want to do any of the following:

- ◆ Destroy a RAID Group on a RAID Group storage system. (You cannot destroy a RAID Group with LUNs bound on it.)
- ◆ Add disks to the LUN. (If the LUN is the *only* LUN in a RAID Group, you can add disks to it by expanding the RAID Group.)
- ◆ Use the LUN's disks in a different LUN or RAID Group.
- ◆ Recreate the LUN with different capacity disks.

In any of these situations, you should make sure that the LUN contains the disk you want. In addition, if the LUN is part of a Storage Group, you must remove it from the Storage Group before you unbind it.

This section describes how to do the following:

- ◆ Determine which disks make up a specific LUN (next).
- ◆ Remove a LUN from the Storage Groups (page 14-23) that contain it when you know which server uses the LUN or which storage system contains the LUN.
- ◆ Unbind a LUN (page 14-11).

To determine which disks make up a LUN

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the LUN you may want to unbind, and then display the disks in the LUN.

To remove LUNs from storage groups

Removing a LUN from a Storage Group makes the LUN inaccessible to the servers connected to the Storage Group.

1. Right-click the icon for the Storage Group from which you want to remove a LUN, and then click **Select LUNs**.
2. In the **Storage Group Properties - LUNs** tab, move the LUNs you want to remove from **Selected LUNs** into **Available LUNs**.
3. Save your changes and close the dialog box.

To unbind a LUN

You cannot unbind a LUN in a Storage Group until you remove the LUN from the group as described in the previous procedure.



CAUTION

Unbinding a LUN destroys any data on it. Before unbinding a LUN, make a backup copy of any data on it that you want to retain. Do not unbind the last LUN owned by an SP connected to a NetWare or Solaris server unless it is absolutely necessary. If you do unbind it, do the following:

NetWare server - Refer to the Release Notice for the NetWare Navisphere Agent for information on how to bind the first LUN.

Solaris server - Edit the Agent configuration file on the servers connected to the SPs that you will use to bind LUNs.

For each server with access to the LUN that you want to unbind, follow the step below for the operating system running on the server:

AIX or HP-UX - Unmount all file systems on the server associated with the LUN, and deactivate and then export the volume group associated with the LUN.

NetWare - Unmount all volumes on all partitions that are associated with the LUN, and then delete these volumes and partitions.

Solaris - Unmount all partitions that are associated with the LUN.

Windows - Stop all processes on the partitions associated with the LUN and delete the partitions.

To Unbind a LUN:

1. Display the icon for the LUN you want to unbind by doing the following:
 - a. From the **Storage** tab of the **Enterprise Storage** dialog box, double-click the icon for the storage system with the LUN you want to unbind.
 - b. Double-click the icon for the SP that owns the LUN.
2. Right-click the icon for the LUN to unbind, and click **Unbind LUN**.
3. After you confirm the operation, the LUN icon disappears from the trees that contained it.

Using the LUN migration feature to copy a LUN

The LUN Migration feature is not available for FC4700-Series storage systems.

The LUN migration feature, included in the Navisphere Manager software and the Navisphere CLI, lets you copy the data in one LUN or metaLUN to another LUN or metaLUN. You might do this to

- ◆ change the type of drive the data is stored on (for example, from more economical ATA to faster FC, or vice-versa)
- ◆ select a RAID type that better matches the data usage, or
- ◆ recreate a LUN with more disk space.

For example, you may have a metaLUN that has been expanded several times by concatenation with other LUNs (not by addition of another entire disk unit). You can use the migration feature to copy the metaLUN onto a new larger LUN, which, being a single entity and not a group of several entities, provides more appropriate performance characteristics.

During a LUN migration, Manager copies the data from the source LUN to a destination LUN. After migration is complete, the destination LUN assumes the identity (World Wide Name and other IDs) of the source LUN, and the source LUN is destroyed.

Using Manager, you can start migrations, display and modify migration properties, and display a summary of all current migrations on one storage system or on all the systems in the domain. You can also cancel a migration, which deletes the destination copy and restores the storage system to its original state.

If the storage system is connected to a VMware ESX Server, and you are migrating the data on a LUN or metaLUN to a larger LUN or metaLUN, any participating LUNs or metaLUNs must be configured as raw device mapping volumes, set to compatibility mode. In order for the ESX Server and virtual machines to recognize the migrated data, rescan the bus at the ESX Server level, and the virtual machine level (optional), after the migration completes.

If the LUNs or metaLUNs are not configured as a raw device mapping volumes, use the VMware **vmkfstools** utility to reconfigure them. For information on using this utility, refer to *Managing Raw Device Mappings* technical note found at

http://www.vmware.com/pdf/esx25_rawdevicemapping.pdf.

LUN migration requirements

- ◆ The source and destination LUNs for the migration *must* reside in the same storage system.
- ◆ You can launch only one LUN migration in a single operation.
- ◆ The source LUN cannot be a hot spare, or a binding, expanding, migrating, or private LUN.
- ◆ The source LUN cannot be a component LUN of a metaLUN.
- ◆ The destination LUN must be the same size or larger than the source LUN.
- ◆ The destination LUN cannot be a hot spare, a binding, expanding, migrating, or private LUN.
- ◆ The destination LUN cannot be in a Storage Group.
- ◆ The destination LUN cannot be part of a SnapView, MirrorView, or cloning operation.

LUN migration limits

The number of supported active and queued migrations is based on the storage system type. Table 14-5 list these limitations.

Table 14-5 LUN migration limits

Storage system type	Active migrations	Queued migrations
CX700	12	1012
CX600	4	508
CX500	8	504
CX400	4	252
CX300	4	252
CX200	4	124

Access to migration objects

Any Navisphere administrator, manager, or monitor can display the LUN migration objects in the Enterprise Storage window. However, only an administrator or manager can start or cancel a LUN migration, or modify its properties.

Migration states

A migration can be in any of the following states:

- ◆ Queued
- ◆ Migrating
- ◆ Faulted
- ◆ Migrated
- ◆ Transitioning

Migrating a LUN

To migrate a LUN:

1. Ensure that the source and destination LUNs meet the requirements described above.
2. In the **Storage** tree, navigate to the LUN or metaLUN whose data you want to migrate (source LUN).
3. Right-click the LUN or metaLUN and click **LUN Migration > Migrate**.
4. Specify a destination LUN from the list of available LUNs.
5. Specify a migration rate: **ASAP**, **High**, **Medium**, or **Low**.
6. If the LUN choices and migration rate you specified are correct, click **OK**.

Manager displays a confirmation box, and, if the destination LUN is larger than the source LUN, a warning. The warning states that you are changing the LUN size and should verify that all attached hosts can handle such a dynamic change.

7. To continue with the migration, select **Yes**.

The LUN migration starts.

Manager displays the estimated time to complete a migration in the **Migration Summary** dialog box, and also on the migration property page. For the first poll cycle, the estimated time to complete is displayed as a question mark (?) because the software needs one poll cycle to calculate a data transfer rate.

When a migration operation ends normally, Manager destroys the source LUN or metaLUN. The newly migrated destination LUN or metaLUN, which has the same identifiers as the original, is ready for use. If the source LUN or metaLUN was part of a Storage Group originally, Manager moves the destination LUN into that Storage Group.

After the migration completes, you can use the newly migrated LUN or metaLUN exactly the same as the original. You can write to it or use it with SnapView or MirrorView. Host connectivity continues, so attached servers can see the LUN or metaLUN and use it without knowing the migration took place.

Migration status - LUN properties

While a LUN or metaLUN is part of an active LUN migration, either as a source or destination, Manager adds a tab, labeled **Migration**, to its property dialog box. This tab appears only while the migration is in process.

Clicking the **LUN Migration Properties** tab displays the

- ◆ Destination LUN ID
- ◆ Destination LUN name
- ◆ Migration rate
- ◆ Migration state
- ◆ Percent complete
- ◆ Estimated time to complete migration

Canceling a migration

To cancel a migration, on the **LUN Migration Properties** tab, click **Cancel**. Manager displays a confirmation box. If you confirm, Manager cancels the migration and deletes any copied material on the destination LUN.

Displaying a migration summary

To display a summary of a storage system's active migrations, right-click the storage system and select **Migration Summary**. Manager displays the **LUN Migration Summary** dialog box for that storage system.

To display a summary of all active migration within the storage domain, in the **Tools** menu, click **Migration Summary**. Manager displays the **LUN Migration Summary** dialog box for all storage systems in the domain.

Reconfiguring RAID groups

After you create a RAID Group, you can change all its properties except its RAID Group ID and, if the Group has a LUN bound on it, the RAID type it supports.

You can add disks to a RAID Group, but if you want to change the disks in a Group to disks with a different physical capacity, you must destroy the RAID Group (thus unbinding all its LUNs), and then recreate it with the other disks.

This section describes how to do the following:

- ◆ Change the expansion/defragmentation priority property or the automatically destroy after last LUN unbound property of a RAID Group (page 14-17)
- ◆ Expand a RAID Group by adding disks to it (page 14-18)
- ◆ Defragment a RAID Group (page 14-19)
- ◆ Destroy a RAID Group (page 14-21)

Changing the expansion/defragmentation priority or the automatically destroy after last LUN unbound property of a RAID group

You can change the expansion/defragmentation priority property of a RAID Group or automatically destroy the RAID Group after the last LUN is unbound without affecting the data on any of its LUNs.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group whose priority you want to change, and select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. In the **General** tab, set the **Expansion/Defragmentation Priority** and the **Automatically Destroy After Last LUN is Bound** settings, and then click **OK** to apply the settings.

Expanding a RAID group

You can expand a RAID Group by adding unbound disks to it. Expanding a RAID Group does *not* automatically increase the user capacity of already bound LUNs. Instead, it distributes the capacity of the LUNs equally across all the disks in the Group, freeing space for additional LUNs.

If you expand a RAID Group that has only *one* bound LUN with a user capacity *equal* to the user capacity of the RAID Group, you can choose to have the user capacity of the LUN equal the user capacity of the expanded Group. Whether you can actually use the increased user capacity of the LUN depends on the operating system running on the servers connected to the storage system.

If the storage system supports metaLUNs, the **Expand LUN with RAID Group** feature is unavailable.

You cannot expand a RAID Group that supports the RAID 1, Disk, or Hot Spare RAID type because a RAID 1 type must have exactly two disks, and the Disk and Hot Spare RAID types must have only one disk. The number of disks you can use for the other RAID types are listed in Table 14-6:

Table 14-6 Number of disks you can use in RAID types

RAID Type	Number of Disks You Can Use
RAID 5 ^a	3 through 16
RAID 3	5 or 9 (FC series only)
RAID 1/0	4, 6, 8, 10, 12, 14, or 16
RAID 0	3 through 16

a.If you are using 181-Gbyte disks, the maximum number of disks you can use is 12 or the size of the RAID Group cannot exceed 2 terabytes.

To Expand a RAID group

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group that you want to expand, right-click, and then select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. In the **Disks** tab, click **Expand** to display the **RAID Group expansion** dialog box.
3. For an FC-series storage system, if the disks you want to add to the RAID Group are in one enclosure, then in the **Select from** list, click that enclosure.

All disks in a RAID Group must have the same capacity to fully use the storage space on the disks. The capacity of a RAID Group that supports the Hot Spare RAID type must be at least as great as the capacity of the largest disk module in any LUN on the storage system.

4. Under **Available Disks**, for each disk that you want to add to the RAID Group, click the icon for the disk and move it into **Selected Disks**.
5. If the RAID Group contains *only* one LUN with a user capacity equal to the RAID Group's user capacity and you want that LUN's user capacity to increase by the user capacity of the added disks, select the **Expand LUN with RAID Group** check box. Otherwise, clear the check box.

If the storage system supports metaLUNs, the **Expand LUN with RAID Group** feature is unavailable.

6. When **Selected Disks** contains only the icons for the disks you want to add to the RAID Group, click **OK** to begin the operation.

What next?

What you do next depends on whether you cleared or selected the **Expand LUN with RAID Group** check box.

Check box cleared - You can bind additional LUNs on the RAID Group.

Check box selected - You need to make the additional space on the LUN available to the operating system. Refer to the documentation for the operating system.

Defragmenting a RAID group

If you unbind and rebind LUNs on a RAID Group, you may create gaps in the contiguous space across the Group's disks (that is, fragment the RAID Group). This leaves you with less space for new LUNs.

You can defragment a RAID Group to compress these gaps and provide more contiguous free space across the disks. This section explains how to determine the amount of usable (contiguous) free space in a RAID Group, and how to defragment a RAID Group.

To determine the amount of usable space in a RAID group

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group you want to defragment, right-click, and then select **Properties**.

For information about the properties in the dialog box, click **Help**.

The amount of unusable free space is **Free Capacity** minus **Largest Contiguous Free Space**. If the **Free Capacity** is at least as large as the total user space on all the LUNs that you want to bind on the RAID Group, then the RAID Group has enough space to bind the LUNs. However, if **Largest Contiguous Free Space** is less than the total, you do not have enough space available for the LUNs, and you need to defragment the RAID Group.

To defragment a RAID group

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group you want to defragment, right-click, and then select **Properties**.
2. Display the **Partitions** tab.

For information on the properties in the dialog box, click **Help**.

3. Click **Defragment**, and then click **Yes** and **OK** in the **Confirmation** dialog box.

Destroying a RAID group

Before you can destroy a RAID Group, you must unbind all the LUNs on it. Unbinding a LUN destroys all the data on it.

Typically, you destroy a RAID Group only if you want to do either of the following:

- ◆ Use its disks in a different RAID Group.
- ◆ Exchange its disks for disks with a different capacity.

In either of these situations, you should make sure that the RAID Group contains the disks that you want.

This section describes how to

- ◆ Determine which disks make up a RAID Group
- ◆ Destroy a RAID Group

To determine which disks make up a RAID group

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the RAID Group you may want to destroy, and then display the disks in the RAID Group.

To destroy a RAID group

Before you can destroy a RAID Group, you must unbind all LUNs on it (refer to "Unbinding a LUN" on page 14-10).



CAUTION

Unbinding a LUN destroys any data on it. Before unbinding a LUN, make a backup copy of any data on it that you want to retain. Do not unbind the last LUN owned by an SP connected to a NetWare or Solaris server unless it is absolutely necessary. If you do unbind it, you will have to do the following:

***NetWare server* - Refer to the Release Notice for the NetWare Navisphere Agent for information on how to bind the first LUN.**

***Solaris server* - Edit the Agent configuration file on the servers connected to the SPs that you will use to bind LUNs.**

For each server with access to any LUN in the RAID Group that you want to destroy, follow the step below for the operating system running on the server:

AIX or HP-UX:

Unmount all file systems on the server associated with each LUN in the RAID Group, and then deactivate and export the volume group associated with each LUN.

NetWare:

Unmount all volumes on all partitions that are associated with each LUN in the RAID Group, and then delete these volumes and partitions.

Solaris:

Unmount all partitions that are associated with each LUN in the RAID Group.

Windows:

Stop all processes on the partitions associated with each LUN in the RAID Group and delete the partitions.

To Destroy a RAID Group:

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the RAID Group that you want to destroy, right-click, and then select **Destroy**.
2. Confirm the operation.

Reconfiguring storage groups

After you create a Storage Group, you can change the following properties:

- ◆ Name
- ◆ LUNs comprising it
- ◆ Servers connected to it

Changing the name of a storage group

You can change the name of a Storage Group using its **Properties** dialog box.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group whose name or sharing state you want to change, right-click and then select **Properties**.

For information on the properties in the dialog box, click **Help**.

2. Make any changes to the name or sharing state, and then click **Apply** to save your changes.

Adding or removing LUNs from storage groups

Removing a LUN from a Storage Group makes the LUN inaccessible to the servers connected to the Storage Group. Adding a LUN to a Storage Group makes the LUN accessible to the servers connected to the Storage Group.

You can add *selected* LUNs to or remove *selected* LUNs from one or more Storage Groups using the **Storage Group Properties - LUNs** tab.

To add or remove LUNs from a storage group

The following procedure assumes that you have already created the LUN that you want to add to the Storage Group:

1. Right-click the icon for the Storage Group to which you want to add LUNs or from which you want to remove LUNs, and click **Select LUNs**.
2. To add LUNs to the Storage Group — In the **Storage Group Properties - LUNs** tab, move the LUNs you want to remove from **Available LUNs** into **Selected LUNs**.
3. To remove LUNs from the Storage Group — In the **Storage Group Properties - LUNs** tab, move the LUNs you want to remove from **Selected LUNs** into **Available LUNs**.
4. Click **OK** to save your changes.

For a description of the properties in the dialog box, click **Help**.

The **Storage** and **Hosts** trees are updated to reflect the change.

Connecting servers to a storage group or disconnecting servers from a storage group

Connecting a server to a Storage Group makes the LUNs in the Storage Group accessible to the server. Disconnecting a server from a Storage Group makes the LUNs in the Storage Group inaccessible to the server.

1. Open the **Storage Group Properties - Hosts** tab in one of the following ways:
 - Right-click the icon for the Storage Group you want to connect to or disconnect from a server, and click **Connect Hosts**.
 - Right-click the host icon for the server that you want to connect to or disconnect from a Storage Group, and click **Connect Storage**.
2. In **Show Hosts**, select which hosts will display in the **Available Hosts** list.
3. For servers that you want to connect to the selected Storage Group, do the following:

In the **Available Hosts** list, select and move the servers into the **Host to be connected** list.

A server can be connected to only one Storage Group. If a server is already connected to a Storage Group when you connect it to the selected Storage Group, it is disconnected from the first Storage Group and, can no longer access LUNs in this group.

4. For the servers that you want to disconnect from the selected Storage Group, do the following:

In the **Hosts to be connected** list, select and move the servers into the **Available Hosts** list.

5. Save your changes and close the dialog box.

The **Storage** and **Hosts** trees are updated to reflect the change.

Destroying storage groups

All servers connected to a Storage Group lose access to the LUNs in the Storage Group after you destroy the group. The LUNs in a Storage Group are not unbound when you destroy it.

To destroy one or more storage groups

From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group you want to destroy, right-click, and then select **Destroy**.

The **Storage** and **Hosts** trees are updated to reflect the change.

After you set up a storage system, you may want to do reconfigure a storage system.

This chapter describes

- ◆ Upgrading a storage system to support caching..... 15-2
- ◆ Replacing disks with higher capacity disks 15-4
- ◆ Connecting a new server to a shared storage system..... 15-6
- ◆ Disconnecting a server from a shared storage system..... 15-6

For security reasons, we recommend that you place all storage systems in a domain before you configure them.

To configure the storage system you must be logged in as global or local administrator, or global or local manager.

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

Upgrading a storage system to support caching

To upgrade a storage system to support caching, the system operator or service person must install the necessary hardware components, and then you must set up array caching. This section describes how to perform each of these tasks.

Installing the hardware components for caching

All storage systems support read caching. A storage system supports write caching only if it has the required hardware, which varies with the storage-system type, as shown in Table 15-1.

Table 15-1 Hardware requirements for write caching

Icon	FC4700, FC4400/4500, FC5600/5700	FC5200/5300	C1900, C2x000, C3x00	C1000
Disks	0-0 through 0-8	0-0 through 0-4	A0, B0, C0, D0, E0	A0 through A4
SPs	Two		Two, with at least 8 Mbytes memory	
Power supplies	Two in DPE and each DAE		Two	
LCCs			Not applicable	
Backup power	Fully charged SPS		Fully charged BBU	

The system operator or service person can install memory modules, disks, a BBU or SPS, and a second SP, LCC, or power supply without powering down the storage system. For an FC series storage system, the DPE, DAE, and SPS installation and service manuals describe how to install an SP, LCC, power supply, and SPS. For a C series storage system, the storage-system installation and service manual describes how to install the SP, power supply, and BBU. If you add disks, you will probably want to create new LUNs with them.

If you add a second SP, you may want it to own some of the LUNs. You can switch the ownership of a LUN from one SP to the new SP (page 14-8).

Setting up caching

1. Assign memory to the partitions for the caches you will use (page 7-10).
2. Enable the storage-system (SP) caches that you will use, and set the other storage-system cache properties (page 7-12).
3. Enable read or write caching for each LUN that you want to use read or write caching, as follows:
 - a. From the **Storage** tab of the **Enterprise Storage** dialog box, double-click the icon for the storage system with the LUN that will use caching.
 - b. Double-click the SP icon that owns the LUN.
 - c. Choose **Properties** from the right-click menu.
 - d. Display the **Cache** tab.
 - e. Select or deselect the **Read Cache Enabled** and/or **Write Cache Enabled** check boxes to enable or disable caching for the LUN.
 - f. Save your changes and close the dialog box.

Replacing disks with higher capacity disks

You can replace any disks in a storage system with higher capacity disks, as long as you do not replace *all* disks that contain the FLARE software database at the same time. These disks are shown in Table 15-2.

Table 15-2 Database disks

Storage-System Type	Database Disk IDs
FC series	0-0, 0-1, 0-2
C3x00 series	A0, B0, C0, A3, A4
C2x00 series	A0, B0, C0, A3
C1900 series	A0, B0, C0, A1
C1000 series	A0, A1, A3, B0

This section describes how to replace a group of disks that either

- ◆ Does *not* include all the database disks.
- ◆ Does include all the database disks.

To replace a group of disks that does not include all the database disks



CAUTION

Do not power off the storage system during the following procedure.

1. For a non-RAID Group storage system, unbind the LUNs whose disks you want to replace (page 14-10).
2. For a RAID Group storage system, destroy the RAID Groups whose disks you want to replace (page 14-21).
3. Replace each disk one at a time; that is, remove the disk and then insert the replacement disk before removing another disk.
4. Bind the disks into the desired LUNs (page 8-21 for a non-RAID Group storage system; page 8-16 for a RAID Group storage system).

If you replaced one or more database disks, the SP copies the FLARE software from another database disk to the replacement disks.

5. Make the LUNs available to the operating system on the storage-system server as described in the server setup manual for the storage system.
6. Restore any data you backed up from the original LUNs.

To Replace a group of disks that does include all the database disks



CAUTION

Do not power off the storage system during the following procedure.

1. For a non-RAID Group storage system, unbind the LUNs whose disks you want to replace (page 14-10).
2. For a RAID Group storage system, destroy the RAID Groups whose disks you want to replace (page 14-21).
3. Replace each disk one at a time, *except for disk 0-0 in a CX-Series or FC-Series storage system or disk A0 in a C-series storage system*. You will replace these after downloading the FLARE software.
4. Download FLARE software to the storage system.
5. Replace disk 0-0 in an CX-Series or FC-series storage system or disk A0 in a C-series storage system.
6. Bind the disks into the desired LUNs; page 8-21) for a non-RAID Group storage system; page 8-16 for a RAID Group storage system).

When you bind the LUNs, the SP copies the FLARE software to disk 00 or A0 from the other database disks.

7. Make the LUNs available to the operating system on the storage-system server as described in the server setup manual for the storage system.
8. Restore any data you backed up from the original LUNs.

Connecting a new server to a shared storage system

The procedure that follows assumes that the new server is physically connected to the storage system through a switch.

1. From the **Storage** tab of the **Enterprise Storage** dialog box, navigate to the icon for the Storage Group that you want to connect to the new server, right-click, and then select **Connect Hosts**.
2. In the **Available Hosts** list, select the new server and move it into the **Hosts to be Connected** list.
3. Save your changes and confirm the connection.

Disconnecting a server from a shared storage system

1. From the **Storage** or **Hosts** tab of the **Enterprise Storage** dialog box, navigate to the icon for the shared storage system you want to disconnect from the server, and click **Connect Hosts**.
2. In the **Hosts to be Connected** list, select the server to disconnect from the storage system, and move it into the **Available Hosts** list.
3. Save your changes and close the dialog box.

If you are physically disconnecting the server from the storage system, you should remove the connection paths from the server HBA ports to the storage system from the storage system's persistent memory. To do so, use the CLI **port** command. Refer to the CLI manual for information on the **port** command. If you do not remove the connection paths, Manager will continue to display them.

Managing iSCSI (Internet SCSI) Storage Systems

This chapter describes how to initialize an iSCSI storage system and set up iSCSI characteristics such as iSCSI HBAs, iSCSI ports, and CHAP authentication. Topics are

- ◆ About iSCSI storage systems in IP networks 16-2
- ◆ Using iSNS (Internet Storage Naming Service) within an iSCSI storage environment..... 16-5
- ◆ Setting up an iSCSI storage system and server for I/O 16-7
- ◆ Configuring storage system iSCSI port network parameters.... 16-9
- ◆ Configuring the iSNS storage environment 16-10
- ◆ Configuring iSCSI server initiator network parameters 16-12
- ◆ Testing iSCSI network connections 16-13
- ◆ Registering server HBAs or NICs 16-13
- ◆ Setting up CHAP credentials on the storage system 16-14
- ◆ Enabling CHAP authentication on the storage system 16-16
- ◆ Managing CHAP credentials..... 16-17
- ◆ Modifying CHAP credentials..... 16-18

This chapter provides a general overview of the topics discussed. For detailed information on these topics, refer to the online help.

About iSCSI storage systems in IP networks

An iSCSI storage system, such as the CX500i, connects directly to an IP network. Servers with either iSCSI HBAs (host bus adapters) or Ethernet connections can access the iSCSI storage system through an IP network, such as a local-area network (LAN).

A LAN is a set of point-to-point connections between nodes, with each node having its own unique IP address. Connections are made through one or more network components such as switches or hubs. Nodes are connected through a LAN by Ethernet CAT 6 (for Gigabit Ethernet LAN) and CAT 5 (for 10/100 LAN) copper cables. Network switches are *not* nodes.

Each node in an iSCSI environment is either an initiator (server) or a target (storage system). Figure 16-1 shows an initiator node and a target node.




Figure 16-1 Nodes - initiator and target

Before a server is able to initiate server I/O to the iSCSI storage system, it must be configured as follows:

- ◆ You have installed one of the following interface cards and relevant drivers
 - Supported iSCSI HBA cards that have a driver and configuration tool (for example, Qlogic), or
 - Gigabit Ethernet Network Interface Cards (NICs) running Microsoft software that provides HBA functionality.

We support 10Mb, 100Mb, and 1000Mb (Gigabit) Ethernet interfaces, but the storage system only supports 1000Mb. If your NIC does not run GigE, then you need to connect to the storage system using a GigE router or switch.

- ◆ You have cabled the storage system properly (refer to the *Setup Guide*, P/N 300-001-924, that shipped with the storage system).
- ◆ You have installed the Navisphere Server Utilities on each server with access to data on the storage system (refer to the *Setup Guide*).
- ◆ You have installed PowerPath software on the servers for multi-pathing (refer to the *Setup Guide*).
- ◆ You have set the Network parameters and security for the SP management ports on the storage system.

The iSCSI interface uses CHAP (Challenge Handshake Authentication Protocol) to protect the storage system's iSCSI ports from unwanted access. CHAP is *optional*, but if your storage system might be accessed from a public IP network, we strongly recommend that you use CHAP security.

CHAP is a method for authenticating iSCSI users (initiators and targets). The iSCSI storage system can use CHAP to authenticate server initiators and initiators can authenticate targets such as the storage system. To use CHAP security, you must configure CHAP credentials for the storage system iSCSI ports and any servers that will access the storage system data.

If you will be using CHAP security, we strongly recommend that you configure it on both the storage system and the server *before* initiating server I/O.

Figure 16-2 illustrates an iSCSI storage system in a storage domain with various servers and illustrates how iSCSI storage systems differ from Fibre Channel storage systems.




Figure 16-2 FC and iSCSI storage systems in an IP network

Using iSNS (Internet Storage Naming Service) within an iSCSI storage environment

iSNS (Internet Storage Naming Service) is supported only on Windows platforms that are part of an iSCSI network configuration.

The iSNS service provides the same function for TCP/IP storage networks as the Simple Name Server (SNS) service in a Fibre Channel fabric — automated discovery, management and configuration of iSCSI devices. It eliminates the need to manually configure each individual storage device with its own list of initiators and targets. Once configured, the iSNS server assumes responsibility for the discovery and management of iSCSI devices.

The iSNS service includes an iSNS server component and iSNS client components. The iSNS server must reside somewhere within the IP storage network, for example, in the switch firmware, or on a host. An iSNS client resides on both the iSCSI storage system and any iSCSI hosts connected to the storage system. When you start the storage system, the iSNS client on the storage system gathers all the storage system's iSCSI port information and stores it locally on the storage system. When you add a server to the storage system's iSNS configuration, Navisphere establishes a connection from the storage system to the iSNS server, and then registers all the stored information on the storage system with the iSNS server.

To use the features of the iSNS service in an iSCSI storage environment,

- ◆ All the iSCSI requirements listed in the section, "About iSCSI storage systems in IP networks" on page 16-2, must be met.
- ◆ An iSNS server must be running on a network to which the storage system has access.
- ◆ The host initiators or HBAs must be able to use the iSNS service.
- ◆ The storage system must support the iSNS service.
- ◆ The iSNS server must be able to communicate with the storage system.

The following illustration represents a sample iSNS configuration. In this configuration, the management and storage networks are separate. The management station is used to configure both the iSNS servers and the storage system (using Navisphere Manager). The iSNS servers are also connected to the storage network so that the clients may query them for target information. The storage system uses either the storage or the management network to register its targets (ports) with the iSNS servers.




Figure 16-3 Sample iSNS storage configuration

Setting up an iSCSI storage system and server for I/O

For detailed information on configuring iSCSI servers and storage systems for I/O, refer to the *Setup Guide*, P/N 300-001-924, that shipped with the storage system.

Before you initiate server I/O to the iSCSI storage system, you *must* complete all of the following steps (configuring CHAP security is optional).

1. Install the *Navisphere Storage System Initialization Utility* and the *Server Support Utility* on each server (refer to the *Setup Guide*).
2. Run the Storage System Initialization Utility on a server on the same subnet as the storage system to set the IP address, gateway, and subnet mask for each SP management port.
3. Set up the storage system's iSCSI port, setting the set IP address, gateway, and subnet mask for each SP iSCSI port. See "Configuring storage system iSCSI port network parameters" on page 16-9.
4. If your IP network is configured with an iSNS service, add the iSNS servers to the iSNS storage environment (refer to "Configuring the iSNS storage environment" on page 16-10). If not skip this step.

Currently EMC supports the iSNS service only for servers running the Windows operating systems.

5. Each server that needs to access an iSCSI storage system must be configured as an initiator to access the storage system. Set up the iSCSI initiators as described in "Configuring iSCSI server initiator network parameters" on page 16-12.
6. Use the Navisphere Server Utility to register any iSCSI HBAs or NICs that will have access to the storage system data (see "Registering server HBAs or NICs" on page 16-13).
7. (Optional) Configure and enable CHAP authentication as follows:
 - a. Configure CHAP security on the storage system using Navisphere Manager. See "Setting up CHAP credentials on the storage system" on page 16-14.
 - b. Configure CHAP security on each iSCSI initiator using the appropriate vendor tools.

If the server is equipped with NIC cards (Network Interface Cards), enable CHAP initiator authentication on the storage system *before* you configure CHAP credentials on the server. If the server is equipped with Qlogic iSCSI HBAs, enable CHAP initiator authentication on the storage system *after* you configure CHAP credentials on the server (see 16-16).

- c. Enable CHAP initiator authentication on the iSCSI storage system using Navisphere Manager (see "Enabling CHAP authentication on the storage system" on page 16-16).
8. If configuring a storage system for the first time
 - a. Use Navisphere Manager to create RAID Groups, create LUNs, and assign LUNs to servers.
 - b. Run OS management tools to prepare the LUNs for server I/O (scan disks, assign drive letters to LUNs).
 - c. Use the Host Agent installed on servers connected to the storage system to assign LUN mapping. You can stop and start the Agent or reboot the server to assign LUN mapping.

For CX-Series storage systems, the Server Utility does *not* assign LUN mapping.

9. Use the **Connectivity Status** dialog box in Navisphere Manager to verify that all of the network connections are present.

Configuring storage system iSCSI port network parameters

Using Navisphere Manager, set the IP address, gateway, and subnet mask for each SP iSCSI port as follows:

Make sure that you record all port configuration information.

1. In the **Storage** tree, right-click the iSCSI storage system and select **iSCSI > Port management**.

This opens the **Port Management** dialog box for the selected storage system.

Each port object is identified by port number, port type, IP address, and name. Port types include iSCSI and FC (Fibre Channel).

2. Select the desired iSCSI port and click **Properties**.
3. In the **Port Properties** dialog box, under TCP/IP Configuration, set the IP address, gateway, and subnet mask for the iSCSI port, where:

IP address	The IP address of the selected iSCSI port.
Gateway	The network gateway address that the selected iSCSI port uses.
Subnet Mask	The network mask for the selected iSCSI port.

4. Click **Apply** to save the settings and return to the **Port Management** dialog box.
5. Repeat this process for each iSCSI SP port that you need to configure.
6. Click **OK** to close the **Port Management** window.

What next?

- ◆ If your IP network is configured with an iSNS service, continue to the next section, "Configuring the iSNS storage environment".
- ◆ If your network is not configured with an iSNS service, complete the iSCSI network configuration by configuring your iSCSI server network parameters (proceed to the section, "Configuring iSCSI server initiator network parameters").

Configuring the iSNS storage environment

Currently EMC supports the iSNS service only on servers running the Windows operating systems.

When you add an iSNS server to the storage system's iSNS environment, a connection between the server and the storage system is established.

The storage system actively communicates with only one iSNS server at a time (the primary server), but you can add additional servers to act as backup servers.

1. In the **Storage** tree, right-click the icon for the storage system for which you want to establish an iSNS communication path, and click **iSCSI > iSNS**.

The **iSNS Servers** dialog box opens.

2. In **iSNS Servers**, enter the IP address of an iSNS server that you want the storage system to communicate with, and click **Add**.

Navisphere Manager displays a confirmation dialog box, asking if you want to continue with the add operation.

3. Click **Yes** to continue with the add operation, or click **No** to cancel it.

If you click **Yes** to continue with the add operation, Navisphere establishes and tests the validity of the connection, registers the storage system's iSCSI port information with the iSNS server, and displays a success or error confirmation dialog box when the operation completes. The error message displays a brief reason for the connection failure.

If the add operation is unable to establish and validate the connection, you can still add the server to the configuration, and then test the connection after you repair the error.

4. Click **OK** to close the confirmation dialog box.
5. Select the server that you want to be the primary server (server that the iSNS service is actively communicating with), and click **Set Primary**.

If there is only one iSNS server listed, that server, by default, is the primary server. There will be no backup servers.

6. To test the validity of a connection from the storage system to the iSNS server, select the server and click **Test Connection**.

You can test only one server connection at a time.

Navisphere Manager displays a success or error confirmation dialog box when the operation completes.

7. Click **OK** to close the confirmation dialog box.
8. To delete an iSNS server, select the server and then click **Delete**.

Navisphere Manager displays a success or error confirmation dialog box when the operation completes.

9. Click **OK** to close the confirmation dialog box.

What next?

Complete the iSCSI network configuration by configuring your iSCSI server network parameters (proceed to the next section, "Configuring iSCSI server initiator network parameters").

Configuring iSCSI server initiator network parameters

The Microsoft iSCSI initiator gives the same name to all NICs in a server. This name identifies the server, not the individual NICs. For this reason, only one NIC in a server is able to login to each SP port on a storage system.

Each server initiator that issues I/O requests to an iSCSI SP port must be configured with the correct network parameters. Use the HBA vendor tools (SAN Surfer - for Qlogic QLA4010 HBAs, or Microsoft's iSCSI Initiator Properties tool for NICs) to configure the iSCSI initiators as follows:

For detailed information on configuring the initiator ports, refer to the *Setup Guide*, P/N 300-001-924, that shipped with the storage system.

1. Assign the IP address, subnet mask, and gateway to the server initiator.
2. Set the storage system iSCSI target parameters (SPs) to which the HBA initiator has access (use the IP addresses of the iSCSI SP ports from "Configuring storage system iSCSI port network parameters" on page 16-9).

If your IP network is configured with an iSNS service, enter the IP address of the iSNS server that you added to the storage environment (refer to "Configuring the iSNS storage environment" on page 16-10). Do *not* enter the IP addresses of the storage system's iSCSI ports.

3. For each Qlogic HBA, enable timestamps, set the **Execution Throttle** to 256, and uncheck **Immediate Data**.

What next? Proceed to the next section, "Testing iSCSI network connections".

Testing iSCSI network connections

Before initiating I/O or configuring CHAP security, we recommend that you use the Ping or Trace Route procedure to test the iSCSI network connections. Ping tests the network connectivity from a selected iSCSI SP port to a specified network destination. Trace Route records the route from the iSCSI SP port to a specified network destination, displaying the IP address of each hop.

1. Using Navisphere Manager, open the **Storage** tree and right-click the icon for an iSCSI storage system and click **iSCSI > Port Management**.
2. Select the port whose network connectivity you want to test and click **Ping** or **Trace Route**.
3. Enter the IP address for the destination port, and use the default values or enter new values for **Packet Size**, **Timeout** (Ping only), **Count** and **Delay(s)**.
4. Click **Start** to begin the connectivity test.
5. Test results are displayed in the **Results** text box.

Registering server HBAs or NICs

In order to have access to an iSCSI storage system, you must register the iSCSI HBAs or NICs with each iSCSI storage system. To do this, run the *Navisphere Server Utility* on each server.

Click **Start > Programs > EMC > Navisphere Server Utility**.

For further information on using the utility, click **Help** in the *Navisphere Server Utility* window.

For CX-Series storage systems, if the Host Agent is installed on connected servers, you can also stop and start the Host Agent or reboot the server.

What next?

Configuring CHAP security is *optional*, but if you will access your storage system from a public network, we strongly recommend that you configure CHAP security (proceed to the next section, "Setting up CHAP credentials on the storage system").

Setting up CHAP credentials on the storage system

Set up CHAP credentials on the storage system as follows:

1. In the **Storage** tree, navigate to the storage system for which you want to configure CHAP credentials.
2. Right click over the storage system and select **iSCSI > CHAP Management**.

Navisphere Manager displays the **iSCSI CHAP Management** dialog box that lists all configured CHAP initiators.

3. To add CHAP credentials for an initiator (these are required to use CHAP security), in the **iSCSI CHAP Management** dialog box, click **Add**.

Navisphere Manager opens the **Add Chap Credentials** dialog box.

4. In **Defined For**, select either **Initiator** or **Target**.

The **Defined For** field determines the direction of authentication. **Target** specifies that the server uses the username-secret to authenticate the storage system. **Initiator** specifies that the storage system uses the username-secret to authenticate the server initiators.

5. To allow any initiator with access to the storage system to log into the storage system using the specified username and secret, select **Allow any initiator to log in via this CHAP user name and secret**.

The initiator name is a fixed name that often begins with **iqn**. You can find the name using the SAN Surfer utility or the Microsoft Initiator utility.

If you are configuring target credentials, **Initiator Name** is dimmed and unavailable.

6. If you are configuring initiator credentials, enter the user name by either
 - a. Entering a name in **User Name**.
 - b. Selecting **Use Initiator Name as User Name**.

Manager copies the initiator name into **User Name**.

7. If you are configuring target credentials, Manager automatically enters the storage system name in **User Name**.
8. In **CHAP Secret**, if you want to specify the secret (password) in hexadecimal, click **Secret Specified in Hex**.

For Hex, the valid range is 24 to 32 hexadecimal characters, upper and lower case. Otherwise, for ASCII, the valid range is 12 to 16 printable characters, upper and lower case.

9. In **Confirm Secret**, re-enter the secret string you specified above.
10. Review the dialog box to verify all the entries that you can.
11. Click **Apply** to add the initiator, username, and secret to the iSCSI database.

If you made an error (for example, Secret and Confirm Secret don't match), Navisphere Manager displays an error message and positions the cursor in the offending field.

12. If you want to specify another initiator or target, return to step 3.

What next?

- ◆ If the attached server is equipped with NIC (Network Interface Cards), enable initiator authentication on the storage system, before you configure CHAP security (go to the next section, "Enabling CHAP authentication on the storage system").
- ◆ If the attached server is equipped with Qlogic iSCSI HBAs, configure initiator CHAP security on the server *before* you enable initiator authentication on the storage system (refer to the *CX500 2-Gigabit iSCSI Disk Processor Enclosure (DPE2) Setup Guide* that shipped with the storage system).

Enabling CHAP authentication on the storage system

CHAP authentication is disabled by default. This feature allows you to enable or disable iSCSI initiator and target authentication at the port level. When initiator authentication is configured and enabled, the storage system will authenticate all initiators that log in using that port. When target authentication is configured and enabled, the servers will authenticate the storage system. Enable CHAP authentication for an iSCSI port as follows:

1. Open Navisphere Manager on the storage system that you are configuring.
2. In the **Storage** tree, navigate to the iSCSI storage system; double-click the icon for the iSCSI storage system and click **Physical**.
3. Double-click the appropriate bus-enclosure, and then right-click the SP to which the iSCSI ports belong.
4. Right-click the appropriate port and select **Properties**.
Manager displays the **iSCSI Port Properties** dialog box
5. To enable CHAP authentication, select **Require initiator authentication**, and click **OK** to save the changes and close the dialog box, or click **Apply** to save the changes without closing the dialog box.

What Next?

- ◆ If the attached server is equipped with NIC (Network Interface Cards), configure initiator CHAP security on the server (refer to the *CX500 2-Gigabit iSCSI Disk Processor Enclosure (DPE2) Setup Guide* that shipped with the storage system).
- ◆ If the attached server is equipped with Qlogic iSCSI HBAs, go the next section, "Managing CHAP credentials" to verify the CHAP credentials for the storage system.

Managing CHAP credentials

The **iSCSI CHAP Management** dialog box lists all the existing CHAP credentials for a storage system, lets you add new users, and modify or delete existing credentials.

To open the **iSCSI CHAP Management** dialog box, in the **Storage** tree, right-click the storage system for which you want to view or configure CHAP security and click **CHAP Management**.

The dialog box lists all CHAP credentials for the selected storage system using the format shown in Table 16-1.

Table 16-1 iSCSI CHAP management window

Initiator Name	Shows the name of the initiator (server). Note that this is a predefined name that usually begins with iqn (for iSCSI Qualified Name). You can determine the name by running the SAN Surfer utility (shipped with Qlogic HBAs) or the Microsoft Initiator utility (for Gigabit Ethernet NIC cards, available from the Microsoft website).
CHAP User Name	Shows the username defined for the initiator or target in the authentication database.
Defined For	Describes the direction of authentication. Initiator means that the storage system uses the user secret (password) to authenticate the server. Target means that the server uses the secret to authenticate the storage system. Defining a username and secret for an initiator is mandatory if you want the iSCSI security mechanism to work. Each initiator (server) must have a corresponding secret in the authentication database so that the storage system (target) can authenticate it. Defining a secret for a target is an optional extra level of security that lets the server verify the identity of the storage system.
Add	Use to add new CHAP credentials for the storage system. Always enabled.
Modify	Enabled when you select one and only one record shown in the table. Use this to modify an initiator or target CHAP secret.
Delete	Enabled when you select one or more records shown in the table. Otherwise, disabled. Use this to delete an initiator or a target name.

Modifying CHAP credentials

The only CHAP credential you can modify is an existing CHAP secret. You cannot modify the initiator name or the username.

1. Right-click the storage system for which you want to modify CHAP security and click **CHAP Management**.
2. In the **iSCSI CHAP Management** dialog box, click **Modify**.

This opens the **Modify CHAP Credentials** dialog box.

3. Define a new CHAP secret.

If you enter the secret in HEX, you must select the **CHAP Secret Specified in HEX** check box.

4. Click **Apply** to save the CHAP secret without closing the dialog box, or click **Cancel** to close the dialog box and return to the **iSCSI Management** dialog box.

If you did not click **Apply**, **Cancel** does not save any new information.

Managing Legacy Storage Systems and NAS Systems

With Navisphere Manager 6.X, you can manage NAS systems or storage systems that do not have the Storage Management Server software installed, such as pre-FC4700 storage systems. You do this through a portal system.

This chapter describes how to designate and configure a portal system:

- ◆ Portal management architecture 17-2
- ◆ Adding and configuring a portal..... 17-5
- ◆ Monitoring portal system status 17-8
- ◆ Adding IP addresses to the agent configuration files..... 17-8

This chapter gives a general overview of the topics discussed. For detailed information on these topics, please refer to the online help.

Portal management architecture

A portal system is a storage system, or a Windows NT, Windows 2000, Windows 2003 (32 bit) or Windows XP server that is configured to manage systems that do not have Storage Management Server software installed (such as FC5300, FC4500 and NAS systems). To configure a portal system, install Storage Management Server software on the system selected to be the portal, and then assign managing Agents and NAS systems to the portal.

A portal system forwards management requests to the SP or Host Agent that is managing the storage system, and also collects information from the SP or Host Agent and relays that information to Manager 6.X for display updates.

Storage system configured as a portal

Configure a storage system as a portal when you want to use an existing storage system in the domain (system is running Storage Management Server software) as the portal rather than adding a supported Windows Server to the domain. This type of portal configuration provides redundancy since either SP is able to communicate with the Host Agents. If one SP fails, the other SP maintains communication with the legacy systems. Manager displays this portal and any systems that the portal is managing in the **Storage** tree in the **Enterprise Storage** window. Do not use a storage system that has more than 20 attached servers as a portal.


Server configured as a portal

Configure a server as a portal when

- ◆ you *do not* have a storage system with Storage Management Server software installed that can be used as a portal, and you want to manage legacy storage systems and centrally monitor storage system events (refer to "Distributed and centralized monitoring" on page 12-5).
- ◆ you *do* have a storage system with Storage Management Server software installed, but you do not want to use this storage system as the portal.

Since the **Storage** tree displays only storage systems, it does not display an icon for any server configured as a portal. It displays the systems that the portal is managing. Use the **Domain Status** dialog box to check the status of each server portal. Refer to "Monitoring portal system status" on page 17-8.

Figure 17-1 shows a sample Manager 6.X environment with a storage management server (an FC4700 storage system) configured as the portal. The FC4700 storage system marked with an asterisk is the portal to the two FC4500 systems without the Storage Management Server software installed.




Legend: - - - - - Management connection for FC4700s

* Management portal to the two FC4500 Storage Systems
Without the Storage Management Server Software

EMC2251

Figure 17-1 Sample storage system portal environment

Figure 17-2 shows a sample Manager 6.X environment with a Windows 2000 server configured as the portal. In this configuration, the portal server is managing the FC4500 and the FC5300 system - neither of which has the Storage Management Server software installed.



EMC2354

Figure 17-2 Sample server portal environment

Adding and configuring a portal

Before you add and configure a portal, a user with administrator privileges must initialize security and set management server settings for any storage systems or servers that you want to use as portal systems. For storage systems configured as portals, refer to the *EMC Navisphere Security Administrator's Guide*, P/N 069001124. For servers configured as portals, refer to the *EMC Navisphere Manager Installation Sheet*.

The **Portal Configuration** dialog box uses a tree structure to show all portal systems, and any legacy storage systems (storage systems without Storage Management Server software) and NAS systems that each portal system is managing. In order to manage legacy storage systems and NAS systems using Manager 6.X, you must do the following:

- ◆ Designate a storage management server, a Windows XP, Windows Server 2003, Windows 2000, or Windows NT server as a portal (refer to next section, "Adding a portal").
- ◆ Assign managing Agents or NAS systems to the portal (go to page 17-6).

A portal system can be the only system in a domain, or it can be in a domain with other systems running Storage Management Server software.

The **Portal Configuration** dialog box does not reflect the state of the storage systems. It is a tool used to manage legacy storage systems by displaying and building the relationships between the portals and the legacy storage systems. Refer to the **Enterprise Storage** tree or **Domain Status** dialog box for system status.

Adding a portal

To add a storage system or server portal, do the following:

1. On the **Tools** menu in the taskbar, open the **Portal Configuration** dialog box.
2. In the **Portal Configuration** dialog box, right-click the portals container icon and then select **Add Portal**.

In **Available Systems**, Manager lists all systems within the domain that are eligible to be portals - systems that are running Storage Management Server software and are not currently acting as portals.

3. In the **Add Portal** dialog box, select the system that you want to be the portal and move this system into the **Selected Systems** list.
4. Click **OK** to add the portal and close the dialog box.

Manager displays an icon for this portal in the **Portal Configuration** dialog box.

5. In the **Portal Configuration** dialog box, click **OK** to save and apply your changes.

Assigning storage systems and NAS systems to portals

Important Before you can configure systems that the portal will manage, you must edit the Host Agent or the SP Agent configuration files to include the IP address for *each* SP in the storage-system portal, or the IP address for the server portal. Refer to "Adding IP addresses to the agent configuration files" on page 17-8.

A portal is able to manage and monitor events for storage systems that do not have Storage Management Server software installed. You do this by assigning managing Agents or NAS systems to the portal. You assign Agents to a portal by scanning a subnet for the Agents that are used to manage the legacy storage systems, or by entering the IP address of the specific Agent that you want to assign. You can assign any Agent to the portal except those Agents running on storage systems that are also running the Storage Management Server software. You can also scan a subnet for NAS systems or enter the IP address for any NAS systems you want to manage. You can then use Navisphere Manager 6.X to manage the selected systems.

A storage system can be a member of only one domain. An FC4700 storage system, without Storage Management Server software installed, can be managed as a legacy storage system through a portal configuration. If you install Storage Management Server software on the FC4700 storage system, you are able to add this system to another domain. Before doing this, remove the FC4700 storage system from the portal configuration using the **Portal Configuration** dialog box.

Scanning subnets for managing agents and NAS devices

To scan a subnet for managing Agents and NAS systems that you want to assign to the portal, do the following:

1. In the **Portal Configuration** dialog box, select the portal system to which you want to assign managing Agents or NAS devices, and select **Add Storage System**.
2. Select either **NAS Devices** or **Agents** as the discovery type for the subnet scan.
3. Enter the addresses for all the subnets you want to scan and then click **Scan**.

The application starts searching the specified subnets for any NAS devices or managing Agents. When it finds one, it places an entry for the device in the **Discovered Systems** list.

4. From the **Discovered Systems** list, select the systems you want to manage and move them into the **Selected Systems** list.
5. Click **OK** to add these systems to the tree in the **Portal Configuration** dialog box.
6. In the **Portal Configuration** dialog box, click **OK** to save and apply your changes.

Manager adds icons for the selected systems to the tree in the **Portal Configuration** dialog box.

Specifying the IP address of managing agents or NAS systems

To assign a specific managing Agent or NAS device to the portal configuration, do the following:

1. In the **Add Storage System** dialog box, select either **NAS Devices** or **Agents** as the discovery type for the subnet scan.
2. Under **Select Systems**, enter the IP address for the specific system you want to assign to the portal, and then move this system into the **Selected Systems** list.
3. Repeat step 2 until you have selected all the systems you want to assign to the portal.
4. Click **OK** to add these systems to the tree in the **Portal Configuration** dialog box.
5. In the **Portal Configuration** dialog box, click **OK** to save and apply your changes.

Manager adds icons for any storage systems acting as portals and all storage systems assigned to the portals to the tree views in the **Enterprise Storage** window.

Monitoring portal system status

Manager 6.X displays all *storage systems* within the domain (those that have Storage Management Server software installed, and those being managed by a portal) in the **Storage** tree in the **Enterprise Storage** window. However, the **Storage** tree does not display an icon for a server portal since the server portal is not a storage system.

Therefore, you cannot determine the status of the server portal using the **Storage** tree.

For example, if you disconnect a server configured as a portal, Manager will change the status of all systems being managed by this portal to **Inaccessible**. To determine if the problem is with the portal, use the **Domain Status** dialog box. This dialog box displays status information for *all* systems within the domain that are running Storage Management Server software. Refer to "Monitoring domain status" on page 13-24.

Adding IP addresses to the agent configuration files

You can use Navisphere Manager 6.X or Navisphere CLI to configure and manage legacy storage systems (storage systems that are not running the Storage Management Server software).

If you want to use Manager 6.X to configure legacy storage systems you must use a portal. In order to configure storage systems through a portal, you must add the IP address for *each* SP in the storage system portal, or the IP address for the server portal to the SP Agent configuration file for any FC4700 storage systems that are not running the Storage Management Server software, and to the Host Agent configuration file for any pre-FC4700 storage systems.

If you will be using Navisphere CLI to configure storage systems and you want to restrict access to these storage systems, you must add a privileged user, other than the IP addresses, to the Host Agent and SP Agent configuration files of the legacy storage systems (refer to "Adding privileged users" on page 6-3, and "Adding privileged users to the SP agent configuration file" on page 7-20).

If the SP Agent or Host Agent configuration files do not include any privileged users, anyone who can log in to the management station can configure the storage system.

Adding IP addresses to the Host Agent configuration file

To enter an IP address into the Host Agent configuration file, you can manually edit the Host Agent configuration file, or you can use Navisphere CLI.

Using a text editor to edit the Host Agent configuration file

1. Open the Host Agent configuration file using a text editor and add the IP address for *each* SP in the storage-system portal, or the IP address for the server portal to the Host Agent configuration file using the following format:

user system@IP address

where

IP address is the IP address for *each* SP in the storage system portal, or the IP address for the server portal.

For example, for a storage-system portal, you might enter IP addresses similar to the following:

user system@10.11.22.333

user system@10.11.22.334

For a server portal, you might enter an IP address similar to the following:

user system@128.222.33.444

2. Save the **/etc/Navisphere/agent.config** file.
3. Do one of the following to incorporate the changes made to the configuration file:

In Navisphere Manager 6.X, in the **Host Properties - Agent** tab, click the **Reload** button. Doing this incorporates changes to the Host Agent configuration file without the need to stop and restart the Host Agent, or

Stop and restart the Host Agent.

Using Navisphere CLI to edit the Host Agent configuration file

1. Add the IP address for *each* SP in the storage-system portal, or the IP address for the server portal to the Host Agent configuration file using the following CLI command:

```
navicli remoteconfig -setconfig -adduser system@ipaddress
```

where

ipaddress is the IP address for each SP in the storage system portal, or the IP address for the server portal.

For example, for a storage-system portal, you might enter the following commands:

```
navicli remoteconfig -setconfig -adduser system@10.11.12.333
```

```
navicli remoteconfig -setconfig -adduser system@10.11.12.334
```

For a server portal, you might enter the following command:

```
navicli remoteconfig -setconfig -adduser system@10.11.12.444
```

2. Stop and restart the Host Agent.

What next?

You can now use Manager to continue editing the Host Agent configuration file. The **Agent** tab in the **Host Properties** dialog box lets you remotely configure a Navisphere Agent — including basic settings, communication channels, and privileged users — on a supported host.

Adding IP addresses to the SP Agent configuration file

Only CX-Series or FC4700-Series storage systems have SP Agent configuration files.

If the SP Agent configuration file does not include any privileged users, anyone who can log in to the management station can configure the storage system. If you do add the IP addresses for the portal to the configuration file, you potentially lock out Navisphere 5.X users who previously had access.

Add the IP address for *each* SP in the storage-system portal or the IP address for the server portal to the SP Agent configuration file using the following CLI command:

```
navicli -h ipaddress remoteconfig -setconfig -adduser system@ipaddress
```

where

ipaddress is the IP address of the SP whose SP Agent configuration file you are editing.

@*ipaddress* is the IP address for *each* SP in the storage system portal, or the IP address for the server portal.

For example, if the IP address of the SP whose SP Agent configuration file you want to change is 10.15.22.33 and the IP address for the portal is 10.14.22.122, enter the following command:

```
navicli -h 10.15.22.33 remoteconfig -setconfig -adduser system@10.14.22.122
```


Using Trees and the Main Window

Manager uses tree structures to show the storage-system environment it is managing. It displays the trees in its Main window. If other Navisphere applications are installed, you may see additional menu options. For information on these options, refer to the online help or the manual for the specific applications.

This chapter describes

- ◆ Introduction to tree views.....A-2
- ◆ Right-click menu options for basic tree icons A-8
- ◆ Main window A-22

Introduction to tree views

Trees show the relationships between the physical and logical components of managed storage systems. Trees are analogous to the hierarchical folder structure of Microsoft Windows Explorer.

A tree is displayed in the selected tab in the open **Enterprise Storage** dialog boxes in the Main window. The **Storage** tree is displayed in the **Storage** tab, the **Hosts** tree is displayed in the **Hosts** tab, and the **Monitors** tree is displayed in the **Monitors** tab as shown on the following pages.

The **Storage** and **Hosts** trees show icons for the logical components of the managed storage systems and servers. The **Storage** tree also shows icons for the physical components of the managed storage systems and servers to which they are connected. The **Storage** tree shows the icons from a storage-system viewpoint, and the **Hosts** tree shows them from a host viewpoint. The **Monitors** tree shows all Event Monitor configurations composed of centralized and distributed monitoring configurations and any existing monitoring templates.

You use the **Storage** tree to control the physical and logical components of the managed storage systems, the **Hosts** tree to control the LUNs and the storage systems to which the hosts connect, and the **Monitors** tree to create and assign event templates to monitored storage systems.

The managed storage systems are the base components in the **Storage** tree. The **Storage** tree displays a storage-system icon for each managed storage system. The managed hosts are the base components for the **Hosts** tree. In an Access Logix configuration, the **Hosts** tree displays a host icon for each server attached to a storage system. The monitor icons are the base components in the **Monitors** tree. The **Monitors** tree displays a monitor icon for each system in the domain as well as any legacy storage systems and Agents.

You can expand and collapse the storage-system or host icons to show or hide icons for their components (such as SP icons, disk icons, LUN icons, RAID Group icons) just as you can expand and collapse the Windows Explorer folder structure. Use the tree icons to perform operations on and display the status and properties of the storage systems, their components, and their host connections.

Figures A-1, A-2, and A-3 show samples of a partially expanded **Storage** tree, **Hosts** tree, and **Monitors** tree.




Figure A-1 Sample partially expanded storage tree




Figure A-2 Sample partially expanded hosts tree




Figure A-3 Sample partially expanded monitors tree

Specifying the managed storage systems to display

You can display specific managed storage systems using **Filter By** and **Filter For** as listed in Table A-1.

Table A-1 Filters for displaying managed storage systems

Filter by	Filter for
All	N/A
Host	Specific hosts by IP address
Fault Condition	Normal Faulted
Subnet	Storage systems by subnet
Storage System Type	C1000 Series C1900 Series C2x00 Series C3x00 Series FC50xx Series FC5000 Series FC5200/5300 Series FC5400/5500 Series FC5600/5700 Series FC4500 Series FC4700 Series CX700 CX600 CX500 CX400 CX300 CX200 CX200LC Unsupported Inaccessible

When you open a Manager session, it displays one **Enterprise Storage** dialog box with a number. For example, the first Enterprise Storage dialog box you open is **Enterprise Storage 1**; the second is **Enterprise Storage 2**. During the session you can open additional **Enterprise Storage** dialog boxes and close them. You might want one dialog box displaying the **Storage** tree and one displaying the **Hosts** tree.

Selecting tree icons

You select icons in a tree in the same way that you select items in any Microsoft Windows applications.

To select a single icon: — Click the icon.

To select multiple icons: — Do either of the following:

- ◆ Press **Shift** while clicking the first icon and last icon to select the first and last icon and all icons between them.
- ◆ Press **Ctrl** while clicking the icons you want to select.

Icon indicators

In the **Storage** and **Hosts** trees, a storage-system icon represents each managed storage system, and in the **Monitors** tree, a storage-system icon represents each managed storage system that is being monitored for events. Each icon consists of an image and a description. The color of the image and the letter it contains reflect the condition of the component as shown in Table A-2.

Table A-2 Storage-system status icons

Icon	Displayed in	Meaning
	Storage, Hosts, and Monitors trees	The storage system and all of its components are working normally. Manager can communicate with the storage system.
	Storage, Hosts, and Monitors trees	Either the storage system or one or more of its components has failed
	Storage, Hosts, and Monitors trees	Either the storage system or one or more of its components is in a transition state.

Table A-2 Storage-system status icons (continued)

	Storage, Hosts, and Monitors trees	The storage system is unmanaged. Manager is not managing the Host Agent or SP Agent for this storage system
	Storage, Hosts, and Monitors trees	The storage system is inaccessible. Manager cannot communicate with the storage system because the management server software is not running.
	Storage, Hosts, and Monitors trees	The storage system is unsupported. Manager does not support this storage system type.

Right-click menu options for basic tree icons

This section describes the icons that appear in the **Storage, Hosts and Monitors** trees.

Use the right-click menus associated with tree icons to perform operations on and display the status and properties of the storage systems, their components, and their host connections.

To display tree icon menus

Select the icon (s) for the component whose menu you want to display, and right-click.

The only icons that do not have right-click menu options are the **Physical** icon in the **Storage** tree and the container icons, for example, the Storage Groups container icon and the RAID Groups container icon.

Table A-3 describes the right-click options available for each icon in the **Storage** tree; Table A-4 describes the right-click options for each icon in the **Hosts** tree; and Table A-5 describes the right-click options for each icon in the **Monitors** tree.

Each table describes the first occurrence of an icon in the tree.

Table A-3 Basic storage tree icons: images and descriptions

Image	Description	Menu option	Use to
	The storage system and all of its components are working normally.	Create RAID Group	Create a RAID Group from selected disks (RAID Group storage systems only).
		Bind LUN	Bind disks into a LUN.
		Create Storage Group	Create Storage Groups on the storage system (shared storage systems only).
		Faults	Display the Fault Status Report for the storage system.
		Disk Summary	Display a summary of the disks in the storage system.
		Software Operations > Software Installation Wizard	Update existing software or install new software on the storage system.
		Software Operations> Software Operation Progress History	Review the status or track the progress of a software installation. (CX-Series and FC4700-Series only)
		Connectivity Status	Display the Connectivity Status dialog box.
		Update Now	Perform a poll of the selected storage system (only supported on storage systems running Navisphere 6.1 software or higher).
		Properties	Display or set the storage-system properties.
		SnapView (only if software is installed)	Refer to SnapView online help or administrator's guide, P/N 069001180.
		MirrorView (only if software is installed)	Refer to MirrorView online help or administrator's guide, P/N 069001161.
		SAN Copy (only if software is installed)	Refer to SAN Copy online help or administrator's guide, P/N 069001188.
		Manage NAS (IP4700 series only)	Open the web-based, network-attached file server (NAS) device management tool.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Container icon for any hosts connected to the storage system.	None	NA
	Server with name <i>hostname</i> connected to the storage system or to a specific Storage Group.	Properties	Display or set the host properties.
		Update Now	Update the status information for this host (only supported on storage systems running Navisphere 6.1 software or higher).
		Faults	Display the Fault Status Report for the host.
	Physical - Container icon for all the hardware components	None	NA
	Enclosure <i>n</i> ; DAE (Disk- Array Enclosure with enclosure address <i>n</i>	Flash LEDs On	Start flashing lights (LEDs) on an enclosure.
		Flash LEDs Off	Stop flashing lights (LEDs) on an enclosure.
	Disks - Container icon for the disks in the storage system.	None	NA

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Disk - For a CX-Series or FC-Series storage system, the disk in the enclosure and slot identified by <i>diskID</i>, which has the format <i>m-n</i> where <i>m</i> is the enclosure number and <i>n</i> is the slot in the enclosure containing the disk. For a C-series storage system, the disk in the slot identified by <i>diskID</i>, which has the format <i>mn</i> where <i>m</i> is the letter (A, B, C, D, or E) of the SCSI bus for the slot and <i>n</i> is the position on the bus containing the disk.	Software installation	Install new versions of disk firmware.
		Properties	Display or set the disk properties.
	Empty disk - <i>diskID</i> . The disk has been removed from its slot	Same as disk	Same as disk.
	Container icon - Fans in the storage system	None	NA
	Enclosure <i>n</i> Fan A - Drive fan pack in enclosure <i>n</i> in a CX-Series or FC-Series storage system. Enclosure 0 FAN A or B - SP fan pack in enclosure 0 in an FC-Series storage system with a DPE. Enclosure SPE FAN A, B, and C - Fans in fan slot A, B and C in a CX-Series storage system.	Properties	Display or set the fan properties.
	Empty Fan - The fan has been removed from the enclosure.	Properties	Display the fan properties.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Container icon - LCCs in an enclosure of a CX-Series or FC-Series storage system.	None	NA
	Enclosure <i>n</i> LCC A Enclosure <i>n</i> LCC B	Properties	Display the fan properties.
	Empty LCC - The LCC has been removed from the enclosure.	Properties	Display the LCC properties.
	Container icon - Power supplies in the enclosure for a CX-Series or FC-Series storage system or in the storage system for a C-Series storage system.	None	NA
	Enclosure <i>n</i> power supply A - Power supply in power supply slot A in enclosure <i>n</i> in a CX-Series or FC-Series storage system. Enclosure <i>n</i> power supply B - Power supply in power supply slot B in enclosure <i>n</i> in a CX-Series or FC-Series storage system. VSC A - Voltage semi-regulated converter (power supply) in power supply slot A in a C-Series storage system. VSC B - Voltage semi-regulated converter (power supply) in power supply slot B in a C-Series storage system.	Properties	Display the power supply properties.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Enclosure 0 - DPE (Disk-Array Processor Enclosure) in any FC-Series storage system except an FC5000 series. SPE (Processor Enclosure) in any CX-Series storage system	Flash LEDs On	Start flashing lights (LEDs) on an enclosure.
		Flash LEDs OFF	Stop flashing lights (LEDs) on an enclosure.
	Container icon - SP A and SP B	None	NA
	SP A or SP B - In a CX-Series or FC-Series storage system, the SP is in the SP A or SP B slot in enclosure 0. In a C-Series storage system, the SP in the SP A or SP B slot in the enclosure.	View Events	Display the event log for the storage processor (SP).
		Reset Statistics Logging	Set statistics for LUNs, disks, and storage-system caching to zero.
		File Transfer Manager	Transfer files from an SP to the local client.
		Resume	Display information stored in the resume prom for the SP.
		SP Collect	Generate diagnostic log files for the selected SP.
		Properties	Display the properties of the SP.
	Empty SP A or SP B	Properties	Display the properties of the SP.
	Container icon - Standby Power Supplies connected to enclosure 0 of a CX-Series or FC-Series storage system that supports write caching. -or- Battery backups in a C-Series storage system that supports write caching.	None	NA

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Enclosure 0 SPS A - SPS connected to SP A in enclosure 0. Enclosure 0 SPS B - SPS connected to SP B in enclosure 0. BBU - BBU in the storage-system enclosure.	Set Test Time	Set the battery test time.
		Properties	Display the status of the SPS or BBU.
	Container icon for PSM LUN (FC4700 storage systems only)	None	NA
	LUN in an FC4700 storage system reserved exclusively for storage-system SPs to store critical information.	Same as LUN	Same as LUN.
	RAID Group Container icon - RAID Groups on the storage system.	Create RAID Group	Create a RAID Group from selected disks (RAID Group storage systems only).
	Individual RAID Group identified by RAIDGroupID in the storage system. <i>RAIDGroupID</i> is the ID assigned when you created the RAID Group. It is a hexadecimal number between 0x00 and 0x1F. <i>RAIDType</i> is Unbound if no LUNs are bound on the Group. Available RAID types are: RAID 5, RAID 3, RAID 1/0, RAID 1, RAID 0, Disk, or Hot Spare. For example, 0x03[RAID 5].	Destroy	Destroy the RAID Group, unbinding all its LUNs.
		Properties	Display the properties of the RAID Group.
		Bind LUN	Bind disks into a LUN on this RAID Group.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	Container icon - Storage Groups in the storage system or accessible from the host. Individual Storage Group in the storage system or accessible from the host. <i>StorageGroupName</i> is the name of the Storage Group.	Create Storage Group	Create Storage Groups on the storage system (shared storage systems only).
		Destroy	Destroy the Storage Group.
		Select LUNs	Add LUNs to the Storage Group.
		Connect Hosts	Connect servers to the Storage Group.
		Properties	Display the properties of the Storage Group.
		SAN Copy Connections	Connect the Storage Group to a SAN Copy storage system SP port.
	Container icon - LUNs in a Storage Group.	None	NA
	RAID 5 LUN in RAID Group or storage system. LUN <i>LUNID</i> [<i>LUNID</i> , RAID 5]; <i>hostnames</i> - <i>devicename</i> <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number. <i>Hostnames</i> is a list of the names of each server connected to the storage system. <i>devicename</i> is the device name for the LUN on those servers. For MirrorView, see Note at end of table.	Unbind LUN	Unbind the LUN, destroying all the data on it and making its disks available for another LUN or RAID Group.
		Trespass	Transfer ownership of the LUN from the current SP to the peer SP.
		Update Host Information	Scan SCSI devices (including storage systems) connected to all servers connected to the storage system. Updates the Navisphere server information based on the results of the scan.
		Add to Storage Groups	Add the LUN to one or more Storage Groups
		Properties	Display the properties of the LUN.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	RAID 5 in RAID Group or storage system (cont)	SnapView (only if software is installed)	Refer to SnapView online help or administrator's guide, P/N 069001180.
		MirrorView/S or MirrorView/S (only if software is installed)	Refer to the online help or administrator's guide. MV/S P/N 069001161 MV/A P/N 300-001-333
		SAN Copy (only if software is installed)	Refer to SAN Copy online help or administrator's guide, P/N 069001188.
		Expand	Expand the user capacity of a LUN.
	RAID 3 LUN in RAID Group or storage system. LUN <i>LUNID</i> [<i>LUNID</i> ,RAID 3] <i>hostnames - devicename</i> <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.
	RAID 1/0 LUN in RAID Group or storage system. LUN <i>LUNID</i> [<i>LUNID</i> , RAID 1/0] <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.
	RAID 1 LUN in RAID Group or storage system. LUN <i>LUNID</i> <i>LUNID</i> ,[RAID 1] <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.
	RAID 0 LUN in RAID Group or storage system. LUN <i>LUNID</i> [<i>LUNID</i> ,RAID 0; <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu Option	Use to
	Individual disk LUN in RAID Group or storage system. LUN <i>LUNID</i> [Disk; <i>mirrorstatus</i>] <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.
	Hot spare in RAID Group or storage system. LUN <i>LUNID</i> [Hot Spare] <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number.	Same as RAID 5	Same as RAID 5.
	Unowned LUNs - LUNs, such as hot spares, that are not owned by either SP.	Same as LUN.	Same as LUN.
	MetaLUN -Type of LUN whose capacity is the combined capacities of all the LUNs that compose it.	Destroy	Destroy the metaLUN and unbind all the LUNs that are part of it.
		Trespass	Transfer ownership of the metaLUN from the current SP to the peer SP.
		Update Host Information	Scan SCSI devices (including storage systems) connected to all servers connected to the storage system. Updates the Navisphere server information based on the results of the scan.
		Add to Storage Group	Add the metaLUN to one or more Storage Groups
		Expand	Expand the user capacity of a metaLUN.
		Properties	Display the properties of the metaLUN.
		SnapView (only if software is installed)	Refer to SnapView online help or administrator's guide, P/N 069001180.
		MirrorView (only if software is installed)	Refer to MirrorView online help or administrator's guide, P/N 069001161.
		SAN Copy (only if software is installed)	Refer to SAN Copy online help or administrator's guide, P/N 069001188.

Table A-3 Basic storage tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	MetaLUN component - Part of a metaLUN. A set of LUNs that are striped together and are independent of other LUNs in the metaLUN.	Add LUNs	Add LUNs only to the <i>last</i> component of a metaLUN. This operation always uses the stripe expansion type.
Note If the storage system has the MirrorView option, mirrorstatus indicates the LUN's remote mirror status, which can be any of the following: Mirrored - LUN is the primary image LUN of a remote mirror. Mirrored/No Secondary Image - Remote mirror does not contain a secondary image. Secondary Copy - LUN is a secondary image LUN for a remote mirror.			


Table A-4 Basic hosts tree icons: images and descriptions

Image	Description	Menu option	Use to
	Container icon for any hosts connected to the storage system. These hosts are not necessarily managed hosts.	None	NA
	Server with name <i>hostname</i> connected to the storage system, or to a specific Storage Group. These hosts are not necessarily managed hosts.	Properties	Display or set the host properties.
		Update Now	Update the status information for this host (only supported on storage systems running Navisphere 6.1 software or higher).
		Faults	Display the Fault Status Report for the host.
		Refresh Initiators	Update initiator records of all Storage Groups connected to the host.
	Container icon - LUNs to which the host can perform I/O.	None	NA
	RAID 5 LUN in RAID Group or storage system. LUN <i>LUNID</i> [<i>LUNID</i>, RAID 5]; <i>hostnames</i> - <i>devicename</i>] <i>LUNID</i> is the ID assigned when you bind the LUN. It is a hexadecimal number. <i>Hostnames</i> is a list of the names of each server connected to the storage system. <i>devicename</i> is the device name for the LUN on those servers. Note See Table A-3, the Storage tree, for the icon images and descriptions for the other LUN RAID types - RAID 3, RAID 1/O, RAID 1, RAID 0, individual disk, Hot Spare, and unowned LUNs.	Unbind LUN	Unbind the LUN, destroying all the data on it and making its disks available for another LUN or RAID Group.
		Trespass	Transfer ownership of the LUN from the current SP to the peer SP.
		Update Host Information	Scan SCSI devices (including storage systems) connected to all servers connected to the storage system. Updates the Navisphere server information based on the results of the scan.
		Add to Storage Groups	Add the LUN to one or more Storage Groups.
		Properties	Display the properties of the LUN.
		SnapView (only if software is installed)	Refer to SnapView online help or administrator's guide, P/N 069001180.
		MirrorView (only if software is installed)	Refer to MirrorView online help or administrator's guide, P/N 069001161.
		SAN Copy (only if software is installed)	Refer to SAN Copy online help or administrator's guide, P/N 069001188.
		Expand	Expand the user capacity of a LUN.

Table A-4 Basic hosts tree icons: images and descriptions (continued)

Image	Description	Menu option	Use to
	MetaLUN -Type of LUN whose capacity is the combined capacities of all the LUNs that compose it.	Same as Storage tree.	Same as Storage tree.
	MetaLUN component - Part of a metaLUN. A set of LUNs that are striped together and are independent of other LUNs in the metaLUN. Note See Storage tree icons for descriptions of right-click menu options.	Same as Storage tree	Same as Storage tree
	Container icon - Storage Groups in the storage system or accessible from the host. Individual Storage Group in the storage system or accessible from the host. <i>StorageGroupName</i> is the name of the Storage Group.	None	NA
		Destroy	Destroy the Storage Group.
		Select LUNs	Add LUNs to the Storage Group.
		Connect Hosts	Connect servers to the Storage Group.
		Properties	Display the properties of the Storage Group.
	Container icon - Storage systems connected to the host.	None	NA
	The storage system connected to the host.	Same as Storage tree.	Same as Storage tree.

Table A-5 Basic monitors tree icons: images and descriptions

Image	Description	Menu option	Use to
	Container icon for any hosts connected to the storage system that are serving as monitoring hosts in a centralized or distributed monitoring environment.	None	NA
	Server with name <i>hostname</i> connected to the storage system, that is serving as the monitoring Agent in a centralized or distributed monitoring environment.	Monitor System	Select a storage system for the host to monitor.
		View Events	View the monitored events for a storage system.
		Monitor Options	Add or remove monitored hosts, select the network timeout and the log file size.
		Select Group Template	Select and apply a group template to the monitoring Agent in a Centralized monitoring environment.
		View Response Log	Display the contents of the Response Log.
		View Message File	Display the contents of the Message File.
	The storage systems being monitored by the monitoring Agent.	Same as Storage tree	Same as Storage tree.
	Container icon for monitoring templates	Create New Template	Create a new template.
		Create new Template Based On...	Create a new template based on an existing template.
		Import Template	Import a template from another host location.
		Export All Templates	Export all templates to another host location.
		Configuration Wizard	Define your monitoring configuration and to create and apply templates.
	Monitoring template	Stop Using	Stop using this template.
		Delete	Delete this template.
		Properties	Display and edit the properties for this template.

Main window




Figure A-4 Main window

The Main window, as shown in Figure A-4, is common to all Navisphere 6.X management applications. The menu bar and toolbar icons and the menu options available when you right-click an icon in a dialog box in the workspace may vary with the applications installed. All other Main window components and functions are identical for all applications. You will perform all the tasks needed to set up a storage system, monitor its operation, and display its properties from the Main window.

Application icon

The Application icon on the left side of the title bar provides overall status of all storage systems managed by the current Manager session. Refer to "Icon indicators" on page A-7 for a description of each status icon.

Menu bar

From the menu bar in the Main window you can display these menus: **File**, **View**, **Tools**, **Window**, and **Help**.

If other Navisphere 6.X applications are installed, you may see additional menus. For information on these menus, refer to the online help index or the manual for the application.

File menu

Option	Use to
Login/Logout	Open the Navisphere Login dialog box, or log out of current session.
Setup Domain*	
Select Master*	Select the master node for the domain.
Configure Domain*	Add systems to or remove them from a domain.
Domain Status	View the status of all systems in the domain.
New Enterprise Storage Window	Open a new Enterprise Storage dialog box.

*Visible only for users logged in as administrators.

View menu

Option	Use to
Toolbar	Show or hide the toolbar.

Tools menu

Option	Use to
Security*	
Change Password*	Assign a new password.
User Management*	Set up user accounts.
Faults	Display a list of any faulted storage systems and their faulted components.
Trespassed LUNs Status	Display the ID and status of all trespassed LUNs.
Failover Setup Wizard	Set up and configure monitoring environments.
Connect Host	Register and scan host systems.
Software Operation Status	Monitor the status of the software installation operation.
Select Language	Select a new language for the user interface and online help (appropriate language pack must be installed).
Portal Configuration	Assign and configure portal storage-systems in order to manage legacy and NAS systems.
Event Monitor	
Create New Template	Create a new template.
Create New Template Based On	Create a new template based on an existing template.
Find and Resolve Template Conflicts	Display a list of any template conflicts.
Import Template	Import existing templates from one domain into another.
Export Template	Export existing templates from one domain into another.
Configuration Wizard	Create templates using the Event Monitor Wizard.
Open Events File	Browse for the events file.
SnapView	Visible only if SnapView software is installed.
SAN Copy	Visible only if SAN Copy software is installed.
MirrorView	Visible only if MirrorView software is installed.

*Visible only for users logged in as administrators.

Window menu

Option	Use to
Cascade	Cascade the open Enterprise Storage dialog boxes.
Tile Vertically	Tile vertically the open Enterprise Storage dialog boxes.
Tile Horizontally	Tile horizontally the open Enterprise Storage dialog boxes.
Close All	Close all Enterprise Storage dialog boxes.

Help menu

Option	Use to
Help Topics	Display the online help topics.
About	Display the version of each installed Navisphere application.

Toolbar

The buttons on the toolbar in the Main window let you perform specific operations on managed storage systems without selecting them from a tree in the **Enterprise Storage** dialog box. To perform operations on individual storage systems, use the menu associated with the storage-system icon. When you position the cursor over a toolbar button, the system displays a brief description of the button. Table A-6 lists and describes the Toolbar icons.

Table A-6 **Toolbar icon image**

Button	Name	Use to
	New Enterprise Storage Window	Open a new Enterprise Storage Window.
	Help	Open the online help.

Status bar

The Status bar displays the username, scope (global/local), and role (administrator, manager or monitor) of the current logged-in user. It also indicates whether Security has been initialized and provides the IP address for the portal storage system.

Workspace

The workspace in the Main window contains the dialog boxes that you use to perform storage-system tasks and any **Enterprise Storage** dialog boxes that you open.

A

- access control enabled property
 - defined 9-2
 - setting 9-2
- Access Logix
 - defined 1-3
 - Storage Groups 9-1
- adding LUNs 10-5
- adding templates to storage systems 12-33
- Agent
 - editing configuration file
 - AIX server 9-9
 - HP-UX server 9-9
 - NetWare server 9-9
 - Solaris server 9-10
 - Host
 - adding privileged users 6-3
 - defined 1-2
 - editing the Host Agent file for portals 17-9
 - managed, defined 1-3
 - SP
 - adding privileged users 7-20
 - defined 1-4
 - editing the SP Agent configuration file for portals 17-11
- ALPA property of SP 7-18
- application icon, function A-23
- authentication, CHAP, see Chapter 3 16-2
- automatic destroy property
 - defined 8-12
 - setting, after creating RAID Group 14-17
- AX-Series storage systems

- dialog box limitations 4-6
- managing with Navisphere Manager, overview 4-2
- managing, prerequisites 4-2
- right-click menu options 4-6
- terminology differences 4-3
- tree structure, sample 4-5

B

- basic storage component icons A-9, A-19, A-21
- battery
 - self-test 7-20
 - see also* BBU (battery back up) or SPS (standby power supply)
- BBU (battery backup unit)
 - failure states 13-14
 - properties, displaying 13-16
 - self-test 7-20

C

- centralized monitoring
 - configuring for a single domain 12-7
 - configuring for multiple domains 12-9
- CHAP authentication, enabling 16-16
- CHAP authentication, see Chapter 3 16-2
- clearing events in events window 12-39, 13-19
- CLI, using to configure storage systems 6-2
- client, defined 1-2
- connectivity status, monitoring 13-21
- constant prefetch property, defined 14-5
- create templates based on existing templates 12-26
- Customer support B-3

Customer Support Center B-7
 customized responses, deleting 12-29

D

data access control for storage system
 defined 9-2
 enabling 9-2
 database directory, defined 1-2
 database disks 3-14
 default owner property
 defined 8-7
 setting after binding LUN 14-8
 defragmenting a RAID Group 14-20
 deleting
 customized responses 12-29
 templates 12-32
 destroying a RAID Group 14-21
 device information in Agent configuration file,
 editing
 AIX server 9-9
 HP-UX server 9-9
 NetWare server 9-9
 Solaris server 9-10
 diagnostics, server related 7-23
 dialog boxes
 Disk Properties 3-14
 Storage System Properties
 Cache tab 7-15
 Data Access tab 9-2
 directory, defined 1-2
 disable size property, defined 14-6
 disk
 disk type, defined 8-3
 disk, individual
 icon for A-17
 RAID type defined 8-3
 disks
 cache vault 13-8
 database 3-14
 failure states 13-5
 faulted, rebuilding on hot spare 13-7
 in LUN 14-11
 in RAID Group 14-21
 increasing capacity of 15-4
 number in RAID types 8-4
 properties, displaying 13-16
 distributed monitoring, configuring 12-5

domain, defined 1-2
 drive fan pack, faulted 13-10

E

element size property, defined 8-6
 E-mail, setting up responses 12-21
 enable auto assign property, defined 8-8
 enable read cache property, defined 8-7
 enable watermark processing property, defined
 7-12
 enable write cache property, defined 8-7
 enforce fair access property, defined 7-16
 Enterprise Storage dialog boxes
 Equipment tab A-2
 Hosts tab A-2
 Storage tab A-2
 Equipment tree A-2
 event logs, viewing 13-17
 Event Monitor
 adding monitored Agents 12-17
 centralized monitoring guidelines 12-5
 clearing events 12-39
 configuring centralized monitoring for a
 single domain 12-7
 configuring centralized monitoring for
 multiple domains 12-9
 configuring distributed monitoring 12-5
 deleting templates 12-32
 distributed monitoring guidelines 12-5
 event code format 12-36
 importing and exporting templates 12-29
 printing events 12-38
 removing monitored Agents 12-18
 saving events to the log file 12-38
 templates
 creating based on existing templates
 12-26
 creating with customized responses
 12-25
 setting up responses 12-21
 terminology 12-4
 testing configurations 12-28
 testing responses 12-27
 understanding SP event codes and event
 severity 12-34
 viewing system event logs 12-36
 viewing template properties 12-31

- viewing the message file 12-38
- viewing the response log 12-37
- Event Monitor response matrix
 - centralized monitoring 12-3
 - distributed host-based monitoring 12-3
 - distributed SP-based monitoring 12-3
- events
 - clearing 12-39, 13-19
 - filtering 13-18
 - opening log file 13-19
 - printing 12-38, 13-19
 - saving to log file 12-38, 13-19
 - tracking to users 13-18
 - understanding event code format 12-34
 - viewing 13-17
 - viewing details 12-37, 13-18
- expanding a RAID Group 14-18
- expansion/defragmentation priority property
 - defined 8-12
 - setting after creating RAID Group 14-17

F

- failover values, setting
 - using the Failover Setup Wizard 5-3
 - using the Group Edit dialog box 5-10
 - when registering an initiator 5-10
- fair access to storage system, defined 7-16
- fan module, faulted 13-11
- fans, faulted
 - drive fan pack 13-10
 - fan module 13-11
 - SP fan pack 13-11
- faults
 - power supply 13-12
 - SPS 13-13
 - storage system 13-4
- FC-Series storage system, defined 1-2
- File menu A-23
- filtering events 13-18
- filters for displaying managed storage systems A-6
- FLARE software
 - committing on a CX-Series or FC4700-Series storage system 3-11
 - reverting to previous on CX-Series or FC4700-Series storage system 3-12

- upgrading on pre-FC4700 storage system 3-15

H

- HBA
 - reconnecting to Storage Groups 13-23
 - registering 13-22
 - removing initiator records 13-22
- Help menu A-25
- high watermark property, defined 7-12
- host
 - updating host connectivity 13-23
- host access status property, setting 7-5
- Host Agent
 - defined 1-2
 - editing the Host Agent configuration file for portals 17-9
- Host, managed, defined 1-3
- Hosts tree A-2
- hot spare
 - defined 8-3
 - disks that cannot be 8-4
 - icon for A-17
 - rebuilding failed disk 13-7
 - type, defined 8-3

I

- icons
 - application A-23
 - basic storage components A-9, A-19, A-21
 - color of A-7
 - letter in A-7
 - power supplies, what to do when orange 13-12
 - RAID Groups A-14
 - SPS, what to do when orange 13-13
 - VSCs A-12
 - what to do when orange 13-12
- idle count property, defined 14-6
- importing templates 12-29
- individual disk
 - defined 8-3
 - icon for A-17
- initiator, iSCSI 16-17
- iSCSI, see Chapter 3 16-2

L

LAN

- introduced 16-2

- language, selecting new 2-5

- legacy Agents, managing 17-2

log file

- opening 13-19

- printing events 12-38, 13-19

- saving events 12-38, 13-19

- SPs, generating 13-20

- logs, viewing event logs 13-17

- low watermark property, defined 7-12

LUN properties

- default owner

- defined 8-7

- setting after binding LUN 14-8

- default values 8-9

- displaying 13-16

- element size, defined 8-6

- enable auto assign, defined 8-8

- enable read cache, defined 8-7

- enable write cache, defined 8-7

- for different RAID types 8-9

- prefetch

- available values 14-7

- default values 14-6

- defined 14-5

- setting 14-8

- rebuild priority, defined 8-6

- verify priority, defined 8-6

LUNs

- maximum number

- in RAID Group 8-16

- in Storage Group 8-16

- in storage system 8-16

LUNs (logical units) 10-5

- default ownership, transferring to another SP 14-8

- defined 8-2

- disks in 14-11

- expanded, *see* <emphasis>metaLUNs

- expanding 10-5

- restrictions 10-5

- results 10-6

- maximum number in RAID Group 8-12

- monitoring trespassed 13-24

- properties, *see* LUN properties

- RAID types, defined 8-2

- rebuilding failed disk on hot spare 13-7

- reconfiguring 14-2

- unbinding 14-10

- unowned, defined A-17

- see also* LUN properties

M

Main window

- application icon A-23

- menu bar A-23

- File menu A-23

- Help menu A-25

- Windows menu A-25

- toolbar A-25

- workspace A-26

- Managed Agent, defined 1-3

- Managed Host, defined 1-3

- managed server, defined 1-3

- managed storage system

- defined 1-3

- filters A-6

- see also* storage system

- management control for a server

- defined 7-3

- enabling 7-5

- Manager, starting a session 2-3

- master node, defined 1-3

- maximum prefetch property, defined 14-6

- memory partitions, assigning memory to 7-10

- menu bar A-23

- File menu A-23

- Help menu A-25

- View menu A-23

- Window menu A-25

- menus on menu bar

- File A-23

- Help A-25

- View A-23

- Window A-25

- message file, viewing 12-38

metaLUNs

- adding LUNs 10-9

- creating, using the wizard 10-12

- defined 10-3

- destroying 10-19

- expanding 10-9

- restrictions 10-9
- results 10-10
- SnapView, MirrorView and SAN Copy LUNs 10-20
- overview 10-3
- properties, displaying and changing 10-17
- status, monitoring 10-19
- using added capacity
 - for SnapView LUNs 10-20
 - LUNs in clone groups 10-21
 - LUNs in mirrors 10-21
 - LUNs in SAN Copy sessions 10-22
- mirrored write cache property, defined 7-13
- monitored Agents
 - adding to monitoring configurations 12-17
 - removing 12-18
- monitoring
 - metaLUN status 10-19
 - SP events 13-17
 - storage system operation 13-2
 - trespassed LUNs status 13-24

N

- NAS systems, managing 17-2
- NAS, defined 1-3
- Navisphere Agent, *see* Agent
- node, defined 1-3
- node, LAN, defined 16-2
- non-RAID Group storage system
 - binding LUNs 8-21
 - defined 1-3

O

- opening event log file 13-19

P

- packages 1-4
- page size property, defined 7-12
- pager
 - setting up advanced features 12-23
 - setting up responses 12-22
- polling
 - host 13-3
 - manual 13-3
 - storage system 13-3
- portal management

- adding SP IP address to SP Agent
 - configuration file 17-11
- adding SP IP addresses to Host Agent
 - configuration file 17-9
- architecture 17-2
- assigning a portal 17-5
- power supplies
 - faulted 13-12
 - icons for, what do when orange 13-12
 - state, displaying 13-16
- prefetch multiplier property, defined 14-5
- prefetch properties
 - available values 14-7
 - default values 14-6
 - defined 14-5
- prefetch size property, defined 14-5
- prefetch type property
 - defined 14-5
 - setting 14-8
- printing events 12-38, 13-19
- privileged users
 - adding to the Host Agent configuration file 6-3
 - adding to the SP Agent configuration file 7-20
- PROM code, included with FLARE software 3-14
- properties
 - access control enabled
 - defined 9-2
 - setting 9-2
 - BBU, displaying 13-16
 - disk, displaying 13-16
 - enable watermark processing, defined 7-12
 - enforce fair access, defined 7-16
 - high watermark, defined 7-12
 - host access status, setting 7-5
 - low watermark, defined 7-12
 - LUN
 - displaying 13-16
 - prefetch
 - defined 14-5
 - setting 14-8
 - mirrored write cache, defined 7-13
 - page size, defined 7-12
 - RAID Group 8-12
 - displaying 13-16
 - setting after creating Group 14-17

- server, displaying 13-15
- SP A read caching, defined 7-13
- SP A statistics logging, defined 7-7
- SP B read caching, defined 7-13
- SP B statistics logging, defined 7-7
- SPs, displaying 13-16
- storage-system
 - data access
 - defined 9-2
 - setting 9-2
 - displaying 13-15
 - general configuration, defined 7-7
 - hosts, defined 7-16
 - management control, defined 7-2, 7-3
- write caching, defined 7-14

R

- RAID 0
 - defined 8-3
 - LUN, icon for A-16
- RAID 1
 - defined 8-2
 - LUN, icon for A-16
- RAID 1/0
 - defined 8-3
 - LUN, icon for A-16
- RAID 3, defined 8-2
- RAID 5, defined 8-2
- RAID Group storage system
 - binding LUNs 8-16
 - creating RAID Groups 8-14
 - see also* storage system
- RAID Groups
 - creating 8-14
 - defragmenting 14-20
 - destroying 14-21
 - disks in 14-21
 - expanding 14-18
 - icons for A-14
 - maximum number of LUNs 8-12, 8-16
 - properties 8-12
 - automatically destroy 8-12
 - automatically destroy after last LUN unbound 14-17
 - default values 8-12
 - displaying 13-16

- expansion/defragmentation priority
 - defined 8-12
 - reconfiguring 14-17
- RAID types
 - defined 8-2
 - LUN properties for 8-9
 - number of disks in 8-4
- read caching, storage-system, setting properties for 7-15
- rebooting SPs 3-15
- rebuild priority property, defined 8-6
- reconfiguring
 - LUNs 14-2
 - RAID Groups 14-17
 - storage system 3-1, 15-1
- Remote Host Agents
 - adding devices 6-5
 - clearing devices 6-5
 - deleting devices 6-5
 - scanning for devices 6-3
- reserved LUN pool
 - allocating LUNs 11-2
 - configuring 11-5
 - displaying properties 11-7
 - estimating size 11-3
 - managing reserved LUNs 11-3
 - monitoring usage 11-7
- response log, viewing 12-37
- responses
 - deleting customized 12-29
 - setting up 12-21
- retention property, defined 14-6
- Roadmap 9-8, 9-9, 9-10

S

- SAN, defined 1-3
- saving events to log file 12-38, 13-19
- secret (iSCSI password) 16-15
- segment multiplier property, defined 14-5
- segment size property, defined 14-5
- server
 - connecting to
 - shared storage system 15-6
 - Storage Groups 14-24
 - defined 1-3
 - disconnecting from

- shared storage system 15-6
 - Storage Group 14-24
- enabling management control to storage system 7-5
- properties, displaying 13-15
- storage application defined 1-4
- Storage Group for, verifying connection to 9-6
- storage management, defined 1-4
- setting up
 - advanced pager features 12-23
 - E-mail responses 12-21
 - pager responses 12-22
 - responses 12-21
 - SNMP responses 12-24
- shared storage system
 - connecting new server 15-6
 - creating Storage Groups on 9-3
 - defined 1-3
 - disconnecting server 15-6
- SNMP, setting up responses 12-24
- software installation
 - monitoring 3-9
 - overview 3-2
- software operations
 - committing 3-11
 - reverting 3-12
- software packages *see* packages
- SP
 - log files, generating and transferring 13-20
- SP (storage processor)
 - port, ALPA address (SCSI ID), setting 7-18
 - properties, ALPA property 7-18
 - properties, displaying 13-16
 - rebooting 3-15
- SP A
 - read caching property, defined 7-13
 - statistics logging property, defined 7-7
- SP Agent
 - adding privileged users 7-20
 - editing the SP Agent configuration file for portals 17-11
- SP B
 - read caching property, defined 7-13
 - statistics logging property, defined 7-7
- SP events, monitoring 13-17
- SP fan pack, faulted 13-11
- SPS (standby power supply)
 - faulted 13-13
 - icons for, what to do when orange 13-13
 - properties, displaying 13-16
 - self-test 7-20
- state
 - BBU failure 13-14
 - disk failure 13-5
 - power supply, displaying 13-16
- status
 - software packages 3-12
 - trespassed LUNs 13-24
- storage application server, defined 1-4
- Storage Groups 9-1
 - connecting servers 9-5, 14-24
 - creating 9-3
 - defined 1-4, 9-2
 - destroying 14-25
 - disconnecting servers from 14-24
 - maximum number of LUNs 8-16
 - server connections to, verifying 9-6
- storage management server, defined 1-4
- storage system
 - adding templates 12-33
 - caching 7-15
 - caching, write hardware requirements for 7-11
 - configuration overview 1-11
 - data access control
 - defined 9-2
 - enabling 9-2
 - description 7-7
 - disks, replacing with higher capacity 15-4
 - fair access to, defined 7-16
 - faults 13-4
 - maximum number of LUNs 8-16
 - non-RAID Group, defined 1-3
 - operation, monitoring 13-2
 - properties
 - configuration access, setting 7-3
 - data access
 - defined 9-2
 - setting 9-2
 - displaying 13-15
 - general configuration, defined 7-7
 - hosts, defined 7-16
 - management control
 - defined 7-2, 7-3
 - setting 7-2

- RAID Group, defined 1-3
- read cache, setting properties for 7-15
- reconfiguring 3-1, 15-1
- serial number defined 7-7
- shared, defined 1-3
- trees for A-2
- unshared, defined 1-4
- upgrading to support caching 15-2
- write cache, setting properties for 7-15
- see also* shared storage system
- storage system operation, monitoring 13-2
 - connectivity status 13-21
 - faulted disks 13-5
 - faulted fan 13-10
 - faulted LCC 13-9
 - faulted power supply 13-12
 - faulted SPS 13-13
 - faulted SPs 13-8
 - faulted VSC 13-12
 - faults 13-4
- Storage System Properties dialog box
 - Cache tab 7-15
 - Data Access tab 9-2
- Storage tree A-2
- system event logs, viewing 12-36

T

- Technical support B-3
- templates
 - adding to storage systems 12-33
 - creating
 - based on existing templates 12-26
 - with customized responses 12-25
 - deleting 12-32
 - importing 12-29
 - importing and exporting 12-29
 - setting up responses 12-21
 - viewing properties 12-31
- trees
 - defined A-2
 - Equipment A-2
 - Hosts A-2
 - selecting icons on A-7
 - Storage A-2

U

- unbinding LUNs 14-10
- unowned LUNs
 - defined A-17
 - icon for A-17
- unshared storage system, defined 1-4
- upgrading FLARE software on pre-FC4700 storage system 3-14

V

- variable prefetch property
 - defined 14-5
 - setting 14-8
- verify priority property, defined 8-6
- View menu options A-23
- viewing
 - event details 12-37, 13-18
 - event logs 13-17
- VMware ESX Server
 - binding LUNs 8-17
 - connecting to Storage Groups 9-5
 - expanding LUNs 10-6
 - metalUN expanding 10-10
 - migrating data to a larger LUN 14-13
 - using Software Installation Wizard 3-6
 - using vmkfstools to reconfigure LUNs 10-6
- VSCs (Voltage Semi-regulated Converters)
 - icons for A-12
 - what do when orange 13-12
 - see also* power supplies

W

- Window menu A-25
- write caching
 - hardware requirements for 7-11
 - storage-system, setting properties for 7-15