

SGI® InfiniteStorage System Manager (ISSM)
Enterprise Edition
Software Concepts Guide

007-5479-001

COPYRIGHT

© 2002-2008 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as “commercial computer software” subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is Silicon Graphics, Inc., 1140 East Arques Avenue, Sunnyvale, CA 94085-4602.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, and IRIX are registered trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

AIX is a trademark of IBM. HP-UX is a trademark of Hewlett-Packard Company. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Solaris is a registered trademark of Sun Microsystems, Inc. Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States and/or other countries. The program product(s) described in this book is a licensed product of LSI Logic Corporation. LSI Logic and SYM are registered trademarks and RAID Manager and SYMlicity are trademarks of LSI Logic Corporation. All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	May 2008 Original publication.

About This Guide

This guide explains the terminology and features of the SGI InfiniteStorage System Manager Enterprise Edition (ISSM-EE) software for SGI InfiniteStorage products. This guide is written primarily for system administrators.

Note: This guide provides the conceptual framework necessary for understanding the storage management software, but it does not contain specific procedures. See the online Help systems for detailed storage management procedures.

To fully understand this document, you should be familiar with basic networking concepts and terminology.

This guide contains the following chapters:

- Chapter 1, “Management Architecture and Software Components”, describes the major hardware components in the ISSM-EE storage management architecture and introduces the storage management software.
- Chapter 2, “Configuring Storage Arrays”, describes various storage array configuration options to maximize data availability and outlines the high-level steps to configure available storage array capacity into volumes and storage partitions.
- Chapter 3, “Remote Volume Mirroring”, describes Storage Manager Remote Volume Mirroring.
- Chapter 4, “The Volume Copy Premium Feature”, describes basic Volume Copy premium feature concepts, and describes how information is copied between volumes within a single storage array and how other premium features can be used in conjunction with the Volume Copy premium feature.

Software and Hardware Terminology

There have been terminology changes between the current version of the storage management software and previous versions. Also, there have been terminology changes in the accompanying controller and drive enclosure documentation. Refer to the table below if you are upgrading storage array hardware and migrating from earlier storage management software versions.

Old Term	New Term
RAID Module	storage array
LUN ^a	volume
drive group	volume group
remaining capacity	free capacity
unassigned capacity	unconfigured capacity
drive module	drive enclosure
controller module	controller enclosure
battery CRU	battery canister
ESM CRU	environmental card canister
fan CRU	fan canister
power supply CRU	power supply canister
LED	indicator light

a. Logical unit number (LUN) is still used in this version of the storage management software. It is a logical address that the host uses to access a particular volume.

Related Publications

This guide is part of a documentation set. Other documentation in the set includes:

- *SGI InfiniteStorage Hardware User's Guides*. These guides provides hardware installation and site requirements, installation instructions, product overviews, and information about component routine maintenance, operation, troubleshooting, and replacement procedures for SGI's InfiniteStorage product line.
- *SGI InfiniteStorage ISSM-EE Administration Guide* for the storage management software specific to your operating environment. This book helps you:
 - Determine the hardware and software required to install the storage management software.
 - Integrate the necessary hardware components into your network.
 - Install the storage management software.
 - Upgrade controller firmware and configuration settings, if necessary.
 - Identify storage management features unique to your operating environment.

When you have completed the tasks in the *SGI InfiniteStorage ISSM-EE Administration Guide*, you will want to refer to these information resources:

- *Enterprise Management window Help*. Use this online Help system to learn more about working with the management domain.
- *Array Management window Help*. Use this online Help system to learn more about managing storage arrays.

Note: The Help systems contain operating information common to all operating environments. Therefore, you will want to refer to the *SGI InfiniteStorage ISSM-EE Administration Guide* for information specific to your operating environment.

Obtaining Publications

To obtain SGI documentation, go to the SGI Technical Publications Library at:

<http://docs.sgi.com>

Conventions

The following conventions are used throughout this document:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined. Italic typeface is also used for book titles.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. Output is shown in nonbold, fixed-space font.
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
 - SGI
 - Technical Publications
 - 1140 East Arques Avenue
 - Sunnyvale, CA 94085-4602

SGI values your comments and will respond to them promptly.

What's New in This Guide

New Features Documented

This update contains the following new features and software enhancements for the SGI InfiniteStorage (IS) 4500, 4000, and 350:

- RAID 6 on IS350 and IS4000
- Full-config RAID 10 and RAID 0 volume groups:
 - IS4500/TP9700 – up to 112+112 RAID 10 or 224-drive RAID 0
 - IS4000 – up to 56+56 RAID 10 or 112-drive RAID 0
- Greater than 2 TB volumes (LUNs)
- More partitions and logins:
 - IS4500/TP9700 increased to 512 partitions and 2,048 logins (host ports). Up from 64 and 512.
 - IS350 and IS4000 increased to 128 partitions and 512 logins (host ports). Up from 64 and 256.
- More mirrors per system:
 - IS4500/TP9700 increased to 128. Up from 64.
 - IS350 and IS4000 increased to 64. Up from 32.
- More snapshots per volume:
 - IS4500/TP9700 increased to 16. Up from 4.
- IS350 and IS4000 increased to 8. Up from 4.

Contents

Record of Revision	iii
About This Guide.	v
Software and Hardware Terminology	vi
Related Publications	vii
Obtaining Publications	vii
Conventions	viii
Reader Comments.	ix
Figures	xxi
Tables	xxv
1. Management Architecture and Software Components	1
Major Hardware Components	2
Physical Components of a Storage Array	3
Logical Components of a Storage Array	3
Storage Partitioning	4
Network Management of Storage Arrays	5
Network Management Connectivity	5
Types of Network Management Connections	6
Directly Managed Storage Arrays	7
Host-Agent Managed Storage Arrays	7
Major Software Components.	8
Client Software Windows	9
Enterprise Management Window	11
Device Tree	13
Device Table	14
Enterprise Management Window Menus	15

Enterprise Management Window Toolbar 16
Array Management Window 17
Array Management Window Tabs 20
Array Management Window Menu Options 21
Array Management Window Toolbar Options 22
2. Configuring Storage Arrays 25
Volumes and Drive Groups 25
Standard Volumes 25
Drive Groups 26
Drive Group Creation 27
Specifying Volume Parameters 27
Protecting Your Data 29
Software Redundancy 29
Use Remote Volume Mirroring 29
Create a Snapshot Volume 30
RAID Levels and Data Redundancy 30
Dynamic RAID-level Migration 31
Hardware Redundancy 34
Protect Data in the Controller Cache Memory 34
Configure Hot Spare Drives 35
Channel Protection 35
Enable a Background Media Scan 36
Storage Array Password Protection 36
Password Failure Reporting and Lockout 37
I/O Data Path Protection 37
Multi-Path Driver with AVT Enabled 38
Multi-Path Driver with AVT Disabled 38

Creating Volumes and Storage Partitions	39
Volume Groups	39
Standard Volumes	39
Specifying Volume Parameters from Free Capacity	40
Specifying Volume Parameters from Unconfigured Capacity	40
Register the Volume with the Operating System	41
Configuring Storage Array Capacity into Volumes	42
Mappings View	43
Storage Partitioning	47
Storage Partitioning Example	47
Major Steps to Create Storage Partitions	49
Heterogeneous Hosts.	52
Heterogeneous Hosts Example.	53
Maintaining Storage Arrays in the Management Domain	55
Monitoring the Health of Storage Arrays	55
Storage Array Status Icons	57
Event Monitor.	58
Alert Notifications	61
Customer Support Alert Notifications.	62
Problem Notification	63
Recovering from Storage Array Problems	65
Recovery Guru Example	66
Tuning Storage Arrays	70
Task Reference	73
3. Remote Volume Mirroring	79
Feature Concepts and Requirements	79
Introduction to Remote Volume Mirroring	80
Primary and Secondary Volumes	80
Mirror Repository Volumes.	81
Mirror Relationships.	83
Data Replication	83
Write Modes	83
Synchronous Mode	84

Write Consistency Mode 85
Asynchronous Write Mode 86
Resynchronization Methods 86
Manual Resynchronization 88
Automatic Resynchronization 89
General Performance Considerations 89
Link Interruptions or Secondary Volume Errors 90
Connectivity and Input/Output 91
Fibre Channel Fabric Configurations 91
Volume Ownership 92
Distance Between Primary and Secondary Storage Array 94
Using Other Premium Features and Remote Volume Mirroring. 94
Storage Partitioning 94
Snapshot Volume 95
Volume Copy 99
Dynamic Volume Expansion 99
Performing Defragment Operations.	100
Hardware and Software Requirements	100
Hardware Requirements.	101
Software Requirements	101
Recommended Configurations	101
Highest Availability Campus Configuration	101
Campus Configuration	104
Intra-Site Configuration	105
Hardware and Software Checklist	106
Hardware Checklist	106
Software Checklist	107
Hardware and Software Installation	108
Pre-Installation Requirements	108
Site Preparation	108
Hardware and Software Requirements	109
Switch Zoning Overview	109
Hardware Installation.	111

Highest Availability Campus Configuration112
Campus Configuration120
Intra-Site Configuration125
Software Installation131
Verifying Installation and Configuration132
Using the Remote Volume Mirroring Premium Feature134
Enabling and Activating the Remote Volume Mirroring Feature134
Determining Remote Volume Mirroring Feature Status.134
Enabling the Remote Volume Mirroring Feature.136
Activating the Remote Volume Mirroring Feature137
Upgrading Mirror Repository Volumes139
Creating Mirror Relationships142
Creating Volumes for Mirror Relationships142
Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard143
Maintaining Mirror Relationships145
Viewing Storage Array Profile145
Changing the Synchronization Settings and the Order Resynchronization Option147
Suspending a Mirror Relationship147
Suspending Mirrors When Operating in Asynchronous Mode.149
Resuming a Mirror Relationship153
Removing Mirror Relationships154
Changing the Write Mode154
Deleting Primary or Secondary Volumes155
Deactivating and Disabling the Remote Volume Mirroring Feature156
Deactivating the Remote Volume Mirroring Feature157
Disabling the Remote Volume Mirroring Feature157
Disaster Recovery158
Planning for Disaster Recovery158
Backing Up Critical Data158
Preparing the Secondary Storage Array159
Ensuring Failover Protection Using Host Clustering Software.159

Disaster Preparation Checklist	160
Recovering from Switch Failures	160
Recovering from Storage System Failures	161
Entire Site Failover Procedure	162
Rebuilding a Damaged Site	163
Re-creating a Mirror Relationship	163
Reversing the Roles of Primary and Secondary Volumes.	164
Changing a Secondary Volume to a Primary Volume	165
Changing a Primary Volume to a Secondary Volume	165
Troubleshooting	167
Testing Communication	167
General Troubleshooting	168
Inappropriate Fibre Channel Fabric Configurations	175
Storage Array Component Failures	176
General Switch Troubleshooting.	178
Remote Volume Mirroring Critical Events	179
Upgrading Controller Firmware.	181
4. The Volume Copy Premium Feature	183
Volume Copy Concepts	183
Introduction to the Volume Copy Premium Feature	183
Source Volumes	184
Target Volumes	184
Creating a Volume Copy	185
Volume Copy and Persistent Reservations	187
Storage Array Performance	187
Interaction of Volume Copy and Other Premium Features	189
Storage Partitioning	189
Snapshot Volume	191
Remote Volume Mirroring	191
Role Reversals.	193
Volume Copy Restrictions	194

Using the Volume Copy Premium Feature195
Enabling the Volume Copy Premium Feature.195
Creating a Volume Copy.197
Selecting or Creating a Target Volume197
Selecting a Source Volume199
Viewing Drives200
Creating a Volume Copy Using the Create Copy Wizard201
SMrepassist Utility for Windows Operating Systems205
Locating Associated Volumes206
Viewing Volume Copy Properties207
Managing Volume Copies with Copy Manager208
Re-Copying a Volume210
Stopping a Volume Copy212
Removing Copy Pairs213
Changing Copy Priority.214
Setting Target Volume Permissions216
Disabling Target Volume Read-Only216
Enabling Target Volume Read-Only217
Volume Copy Scenarios218
Copying Volumes for Improved Access218
Restoring Snapshot Volume Data to the Base Volume224
Volume Backup and Restoration233
Creating a Backup Copy233
Recovering from a Backup Copy240
Troubleshooting242
General Troubleshooting.242
Volume Copy Critical Events246
Glossary of Volume Copy Terms247
Index253

Figures

Figure 1-1	I/O Data Connection	5
Figure 1-2	Network Management Connections	6
Figure 1-3	Enterprise Management Window	9
Figure 1-4	Array Management Window	10
Figure 1-5	Enterprise Management Window	12
Figure 1-6	Device Tree with a Management Domain	13
Figure 1-7	Array Management Window	19
Figure 2-1	Mappings View Window	43
Figure 2-2	Storage Partitioning Example.	49
Figure 2-3	Defining Hosts and Host Groups in the Topology View	52
Figure 2-4	Host Port Definitions Dialog	53
Figure 2-5	Heterogeneous Hosts Example	54
Figure 2-6	Monitoring the Health of Storage Arrays Using the Enterprise Management Window	56
Figure 2-7	Event Monitor Configuration	59
Figure 2-8	Event Monitor Example	60
Figure 2-9	Problem Notification in the Array Management Window.	64
Figure 2-10	Recovery Guru Toolbar Button	65
Figure 2-11	Recovery Guru Window	67
Figure 2-12	Status Changes During Example Recovery Operation	68
Figure 2-13	Status Changes When an Example Recovery Operation Completes	69
Figure 3-1	Primary and Secondary Volumes Displayed in the Array Management Window	81
Figure 3-2	Mirror Repository Volumes Displayed in the Array Management Window	82
Figure 3-3	Data Replication between Storage Arrays	84
Figure 3-4	Mirrored Volume Pair in Unsynchronized Status	87

Figure 3-5	Mirrored Volume Pair in Synchronization-in-Progress Status	87
Figure 3-6	Controller A as Controller Owner.	92
Figure 3-7	Controller B as Controller Owner	93
Figure 3-8	Storage Partitioning Example	95
Figure 3-9	Snapshot Volume Included in a Mirror Relationship	96
Figure 3-10	Snapshot Volume Associated with Primary Volume Accounting	97
Figure 3-11	Snapshot Volume Failure Because of Role Reversal.	98
Figure 3-12	Highest Availability Campus Configuration.	103
Figure 3-13	Campus Configuration	104
Figure 3-14	Intra-Site Configuration	105
Figure 3-15	Switch Zoning in Cascaded and Non-Cascaded Fibre Channel Switches.	110
Figure 3-16	Highest Availability Campus Configuration.	113
Figure 3-17	Switch Zoning for the Highest Availability Campus Configuration	115
Figure 3-18	Host Bus Adapter Connections to Fibre Channel Switches	117
Figure 3-19	Storage Array Connections to Fibre Channel Switches	118
Figure 3-20	Connecting Switches over a Remote Network to Complete Fabric Environments	119
Figure 3-21	Campus Configurations	120
Figure 3-22	Switch Zoning for the Campus Configuration	122
Figure 3-23	Host Bus Adapter Connections to Fibre Channel Switches	123
Figure 3-24	Storage Array Connections to Fibre Channel Switches	124
Figure 3-25	Connecting Switches over a Remote Network to Complete the Fabric Environment.	125
Figure 3-26	Intra-Site Configuration	126
Figure 3-27	Switch Zoning for the Intra-Site Configuration	127
Figure 3-28	Primary Host Bus Adapter Connections to Fibre Channel Switches	129
Figure 3-29	Secondary Host Bus Adapter Connections to Fibre Channel Switches.	129
Figure 3-30	Primary Storage Array Connections to Fibre Channel Switches	130
Figure 3-31	Secondary Storage Array Connections to Fibre Channel Switches	131
Figure 3-32	Remote Volume Mirroring Feature Status	135
Figure 3-33	Synchronizing (Left) and Optimal (Right) Mirror Relationship	144
Figure 4-1	Creating a Volume Copy	186

Figure 4-2	Select Target Volume and Copy Priority Dialog188
Figure 4-3	Change Copy Priority Dialog189
Figure 4-4	Storage Partitioning Example.190
Figure 4-5	Primary Volume Included in a Volume Copy193
Figure 4-6	Enabling the Volume Copy Premium Feature196
Figure 4-7	Renaming the Target Volume.199
Figure 4-8	Selecting a Source Volume200
Figure 4-9	Examples of the View Drives Dialog.201
Figure 4-10	Create Copy Wizard - Introduction Dialog202
Figure 4-11	Create Copy - Select Target Volume and Copy Priority Dialog203
Figure 4-12	Create Copy Wizard - Preview Dialog204
Figure 4-13	Copy Started Dialog205
Figure 4-14	Volume Properties Dialog.207
Figure 4-15	Copy Manager Dialog209
Figure 4-16	Re-Copy Dialog.211
Figure 4-17	Re-Copying a Volume212
Figure 4-18	Stopping a Volume Copy213
Figure 4-19	Change Copy Priority Dialog215
Figure 4-20	Disabling the Target Volume's Read-Only Attribute217
Figure 4-21	Storage Array Midwest Example.219
Figure 4-22	Creating the Target Volume Finance-Copy220
Figure 4-23	Creating a Volume Copy of Finance221
Figure 4-24	Viewing the Volume Copy Status in the Copy Manager222
Figure 4-25	Defining an Additional Volume-to-LUN Mapping for the Target Volume223
Figure 4-26	Storage Array Midwest Example.225
Figure 4-27	Creating Target Volume Finance-Target.226
Figure 4-28	Creating a Volume Copy227
Figure 4-29	Viewing the Volume Copy Status in the Copy Manager228
Figure 4-30	Disabling the Snapshot Volume229
Figure 4-31	Creating a Volume Copy of Finance-Target.230
Figure 4-32	Viewing the Volume Copy Status in the Copy Manager231

Figure 4-33	Defining an Additional Volume-to-LUN Mapping for the Base Volume	232
Figure 4-34	Storage Array Midwest Example	234
Figure 4-35	Creating a Snapshot Volume of Finance	235
Figure 4-36	Creating the Target Volume Finance-Target	236
Figure 4-37	Creating a Volume Copy the Snapshot Volume Finance-Copy	237
Figure 4-38	Viewing the Status of the Volume Copy in the Copy Manager	238
Figure 4-39	Disabling Snapshot Volume Finance-Source.	239

Tables

Table 1-1	Major Hardware Components	2
Table 1-2	Physical Components of a Storage Array	3
Table 1-3	Logical Components of a Storage Array	3
Table 1-4	Major Software Components	8
Table 1-5	Enterprise Management Window Menus	15
Table 1-6	Enterprise Management Window Toolbar Buttons.	16
Table 1-7	Array Management Window Tabs	20
Table 1-8	Array Management Window Menus.	21
Table 1-9	Array Management Window Toolbar Buttons	23
Table 2-1	Volume Parameters for Free Capacity, Unconfigured Capacity, and Unassigned Drives	28
Table 2-2	RAID Level Configurations	32
Table 2-3	Creating Volumes from Unconfigured or Free Capacity	42
Table 2-4	Mappings View Tab	44
Table 2-5	Volume-to-LUN Terminology	44
Table 2-6	Storage Array Status Indications	57
Table 2-7	Tuning Options in Array Management Window	70
Table 2-8	Task Reference: Populating a Management Domain	73
Table 2-9	Task Reference: Storage Array Configuration	74
Table 2-10	Task Reference: Performance and Tuning	75
Table 2-11	Task Reference: Event Notification	76
Table 2-12	Task Reference: Recovering From Problems	76
Table 2-13	Task Reference: Miscellaneous System Administration	77
Table 3-1	Hardware Checklist	106
Table 3-2	Software Checklist	107
Table 3-3	Variable Definitions	151
Table 3-4	Disaster Preparation Checklist	160

Management Architecture and Software Components

This chapter describes the major hardware components in the InfiniteStorage System Manager Enterprise Edition (ISSM-EE) architecture and the physical and logical components of a storage array. This chapter also introduces the storage management software and describes how storage arrays are managed using this software.

The storage management software may be installed on personal computers or workstations, called storage management stations, anywhere on the corporate network. You send storage management commands from these storage management stations over an Ethernet network to the controllers in the storage array. The commands are either sent directly to the controllers through Ethernet connections on the back of the controllers (directly managed method) or through an Ethernet connection on the host (host-agent managed method). See “Types of Network Management Connections” on page 6 for more details.

The controller firmware executes the storage management commands, and you use the storage management software (which is called the client software) to manage the storage arrays.

Note: You can install the client software on either a host (server) or workstation platform. The storage commands are sent over the host I/O data connection when connected in-band, or through TCP/IP over the network when connected directly (out-of-band).

Major Hardware Components

Table 1-1 lists the major hardware components that are involved in the ISSM-EE storage management architecture. For a detailed description of the network setup requirements, see the appropriate hardware guides for each component.

Table 1-1 Major Hardware Components

Hardware Component	Description
Storage array	Managed by the storage management software and consists of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as volume groups and volumes). Please review the README file and release notes on the CD-ROM for information on supported storage arrays.
Storage management station	Computer you use to manage storage arrays on the network.
Network management station (optional)	Console with installed SNMP-compliant network management software that receives and processes information about managed network devices using SNMP (Simple Network Messaging Protocol).
BOOTP or BOOTP-compliant DHCP server	Used to assign static network-specific information such as an IP address and host name for each controller. This server is not required if you are going to manage all of the storage arrays through the I/O connection using a host-agent. For more information, see “Types of Network Management Connections” on page 6.
Host	Computer, running one or more applications, that accesses the storage array over a fibre channel I/O data connection.

Physical Components of a Storage Array

The physical components that are typically found in a storage array are described in Table 1-2.

Table 1-2 Physical Components of a Storage Array

Component	Description
Controller	Controller board and firmware that controls and monitors the drives and implements the storage management functions.
Controller enclosure	One or more controllers, power supply canisters, fan canisters, and other supporting components in a single enclosure.
Drive	Electro-magnetic mechanical device providing the physical data storage media.
Drive enclosure	Set of drives, power supply canisters, fan canisters, environmental card canisters, and other supporting components in a single enclosure.

Logical Components of a Storage Array

The drives in the storage array provide the physical storage capacity for application data. Use the storage management software to configure the physical capacity into logical components (volume groups and volumes) as described in Table 1-3.

Table 1-3 Logical Components of a Storage Array

Component	Description
Volume group 	Set of drives that are logically grouped together by the controllers in a storage array. Each volume group is assigned an appropriate RAID level to meet application needs for data availability and I/O performance. The RAID level is a volume group parameter, not a volume parameter.
Volume 	Logical object that is the basic structure you create on the storage array to store data. One or more volumes are configured across a volume group, and each volume is seen by the operating system as one drive.

Table 1-3 Logical Components of a Storage Array (**continued**)

Component	Description
Free capacity 	Contiguous region of unused capacity on a designated volume group. The space can be used to create one or more volumes. Note: In the Logical view of the Array Management window, free capacity is displayed as Free Capacity nodes (see Figure 1-7 on page 19). Multiple Free Capacity nodes can exist on a volume group.
Unconfigured capacity 	Capacity that is present in the storage array from drives that have not been assigned to a volume group. The space can be used to create new volume groups and volumes. Note: In the Logical view of the Array Management window, unconfigured capacity is shown as an Unconfigured Capacity node.

Storage Partitioning

Storage partitioning allows you to configure a single storage array as multiple virtual storage arrays up to the maximum value enabled. A *storage partition* is a logical grouping of one or more storage array volumes. Access to a storage partition can be restricted to particular hosts or groups of hosts by defining a set of volume-to-LUN mappings. A *volume-to-LUN mapping* allows you to define what host or host group will have access to a particular volume in the storage array.

For more information on creating storage partitions, see “Storage Partitioning” on page 47.

Network Management of Storage Arrays

Storage array management activities include:

- Setting up a *management domain*, which is a collection of storage arrays to manage.
- Configuring destinations to receive alert messages for critical problems concerning one or more storage arrays in a management domain.
- Configuring available storage array capacity into logical volumes to maximize data availability and optimize application performance.
- Granting various hosts in your storage enterprise access to storage partitions (defined in “Storage Partitioning” on page 47).
- Monitoring storage arrays in the management domain for problems or conditions that require attention.
- Recovering from storage array problems to maximize data availability.
- Tuning for optimal application performance.

To perform these management activities, you use the client software on the storage management station or host to instruct the controllers in the storage array. These commands are sent over the network. The following section describes network management connectivity.

Network Management Connectivity

A storage array receives data from the application host over fibre channel *I/O data connections* (also called I/O data paths), shown in Figure 1-1.

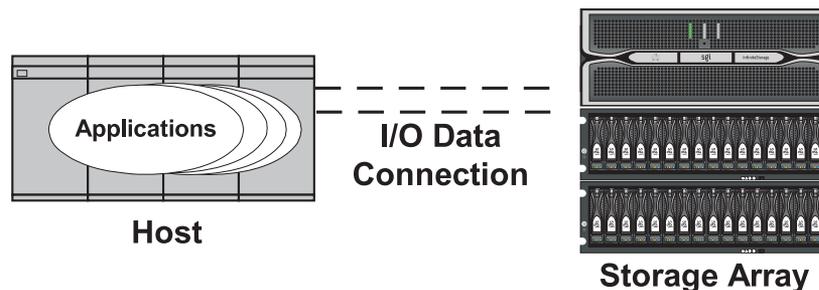


Figure 1-1 I/O Data Connection

Storage management commands are sent from storage management stations over an Ethernet network to the storage array. The connection between the storage management station and the storage array controller, shown in Figure 1-2, is called the *network management connection*. Two kinds of network management connections are described in “Types of Network Management Connections” on page 6.

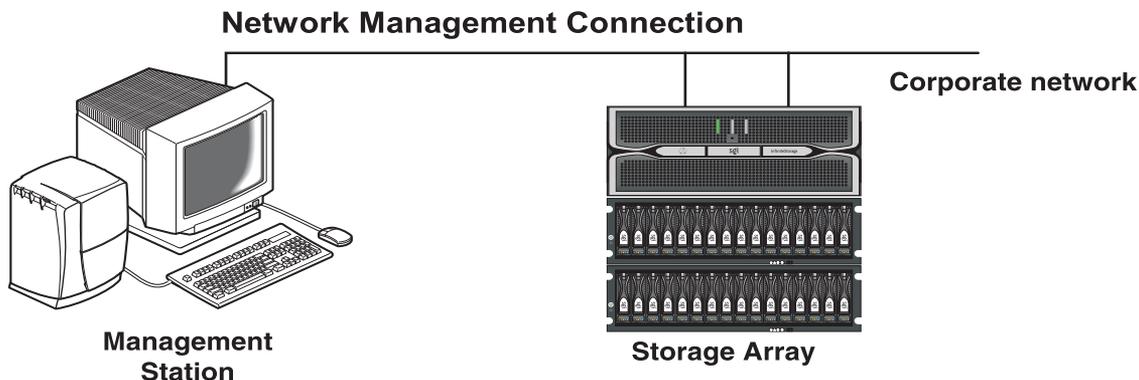


Figure 1-2 Network Management Connections

Types of Network Management Connections

The software allows two different types of network management connections to a storage array: direct or host-agent.

For increased connectivity, you can manage a storage array using a variety of network management connections (direct only, host-agent only, or multiple host-agents). A direct and host-agent combination is also allowed, but is not recommended because the result would be three Ethernet connections (one to each controller and one through the host). When you configure multiple network management connections to a storage array, the storage management software is aware of each connection. When you open an Array Management window and you have multiple connections configured, the software automatically chooses a connection. If a particular connection does not respond, the software tries all other configured network management connections to that storage array.

Directly Managed Storage Arrays

The directly managed method allows you to send storage management commands to the controllers in the storage array directly over the network through each controller's Ethernet connection. To manage the storage array through these Ethernet connections, attach cables from the Ethernet connections on the storage array to the network, define each controller's IP address (or optionally, each controller's host name), and then use the Enterprise Management window to add the storage array to the management domain. This form of RAID management is also referred to as “out-of-band” management.

For more information about directly managed storage arrays, see the *SGI InfiniteStorage ISSM-EE Administration Guide*.

Host-Agent Managed Storage Arrays

The host-agent managed method allows you to send storage management commands through the network connection to a host rather than through the Ethernet connections to each controller. A ISSM-EE storage management software component installed on the host, called a host-agent, receives the requests from the storage management station and sends them to the controllers in the storage array through the fibre channel I/O connection. The communication between the storage management station and the controllers is done through a special volume configured by controllers called an *access volume*. This volume uses one of the available logical unit numbers (LUNs), typically ID 31. See Table 1-4 on page 8 for a description of the host-agent software. This form of RAID management is also referred to as “in-band” management

After you install the host-agent software on the host, add the host to the management domain in order to also add any storage arrays attached to that host. For more information about populating a management domain, see “Device Tree” on page 13.

For more information about host-agent managed storage arrays, see the *SGI InfiniteStorage ISSM-EE Administration Guide*.

Major Software Components

The ISSM-EE storage management software consists of three major components, described in Table 1-4.

Table 1-4 Major Software Components

Component	Description
SMEE Client	<p>You can install the client software in one of two locations: (1) on a storage management station, or (2) on a host. See the <i>SGI InfiniteStorage ISSM-EE Administration Guide</i> for details.</p> <p>For more information, see “Client Software Windows” on page 9</p>
SMEE Agent	<p>The host-agent software can be installed on one or more hosts connected to the storage arrays. The host-agent, along with the network connection on the host, provides a network management connection to the storage array. Instead of using the individual Ethernet connections on each controller, the storage management station can communicate with a storage array through the host that has host-agent management software installed.</p> <p>The host-agent receives requests from the storage management station through the network connection to the host and sends them to the controllers in the storage array through the Fibre Channel I/O connection. The communication between the storage management station and the controllers requires a special volume configured by the controllers called an <i>access volume</i>. This volume uses one of the logical unit numbers (LUNs) available on the host.</p> <p>It is recommended that you install the host-agent software on all host computers, even if you intend to manage the storage array directly over the network. You can stop the host-agent software from running by using an operating system-specific method. See the <i>SGI InfiniteStorage ISSM-EE Administration Guide</i> for details.</p>
SMEE Util	<p>The utility package software contains the utilities that correlates the volumes you create using the storage management software with their operating system device names. See the <i>SGI InfiniteStorage ISSM-EE Administration Guide</i> for details.</p>

Client Software Windows

The client software has two main windows: the **Enterprise Management** window, shown in Figure 1-3, and the **Array Management** window, shown in Figure 1-4.

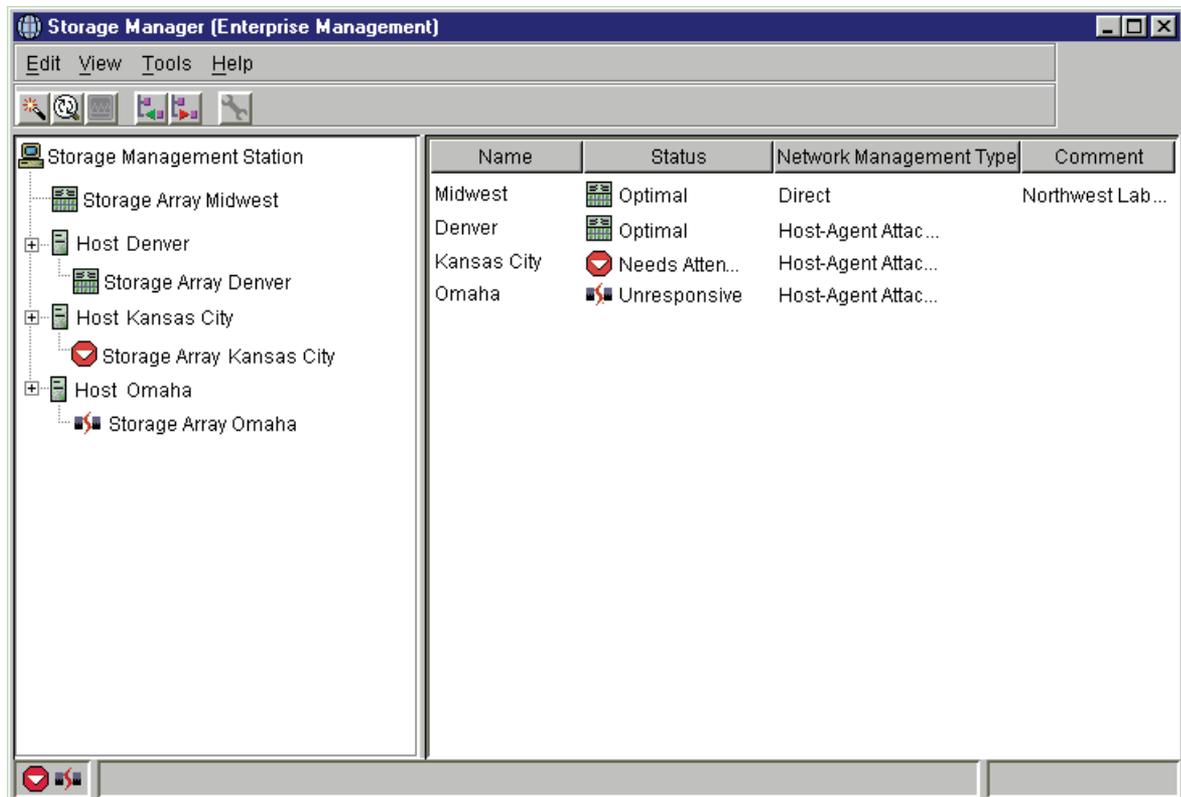


Figure 1-3 Enterprise Management Window

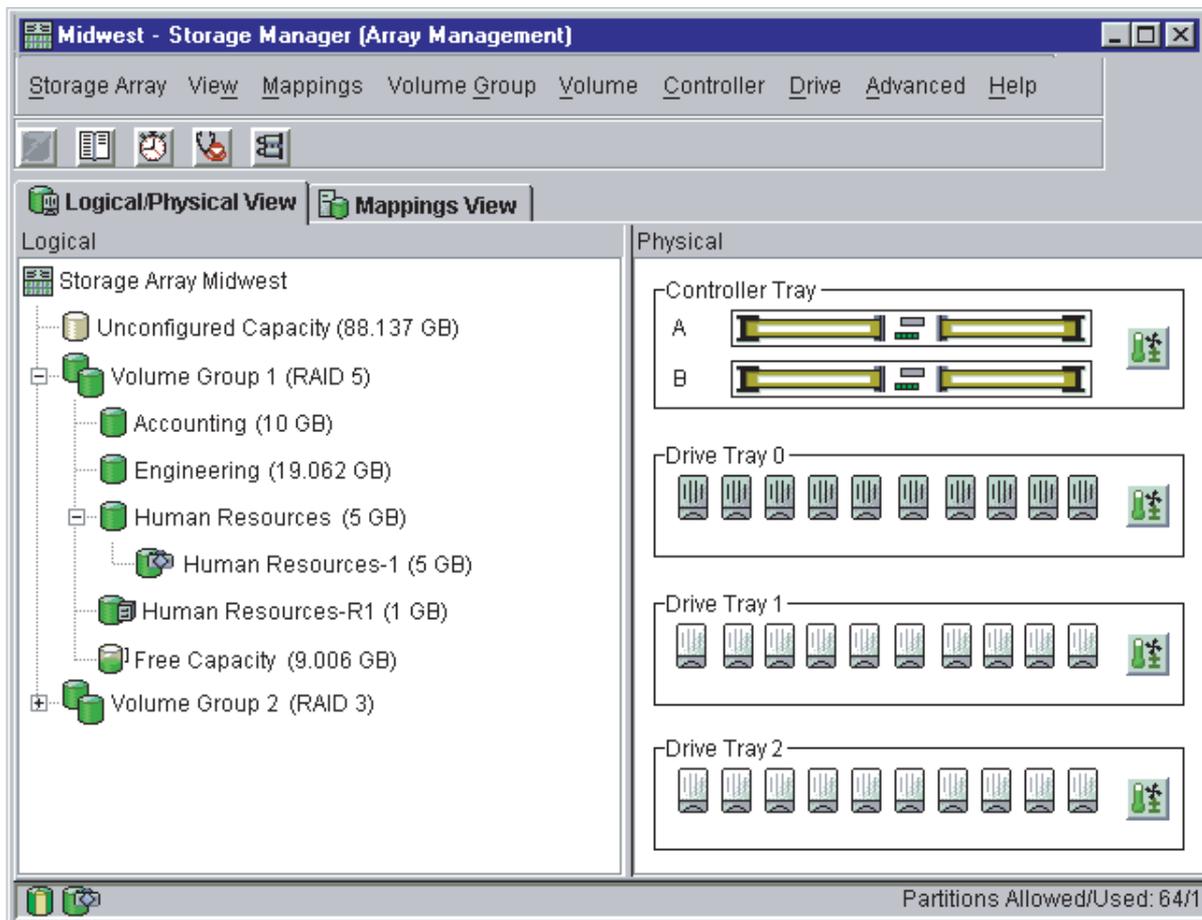


Figure 1-4 Array Management Window

Enterprise Management Window

The **Enterprise Management** window, shown in Figure 1-5, is the first window to appear when you start the software. You use the **Enterprise Management** window to:

- Add and discover the storage arrays you want to manage (see “Related Publications” on page vii).
- Provide a comprehensive view of all storage arrays in the management domain.
- Perform batch storage array management tasks using the Script Editor.
- Configure destinations (e-mail and/or SNMP traps) to receive alert notifications for non-optimal storage arrays.

Note: A local configuration file named emwdata.bin stores all information about the management domain, including the storage arrays you have added and the e-mail or SNMP destinations you have configured. This user-specific configuration file allows each system administrator to create his or her own unique view of the devices in the management domain, without affecting the information shown to other system administrators.

To receive critical alert notifications, the **Enterprise Management** window must be running (it can be minimized).

After adding storage arrays, you use the **Enterprise Management** window primarily to monitor the storage arrays in the management domain for the Needs Attention status and for alert notification of critical errors affecting the storage arrays. When you are notified of a non-optimal storage array status in the **Enterprise Management** window, you can start an **Array Management** window specific to the affected storage array to learn more detailed information about the storage array condition. You can also use the **Array Management** window to perform recovery or other management tasks. For more information, see “Monitoring the Health of Storage Arrays” on page 55.

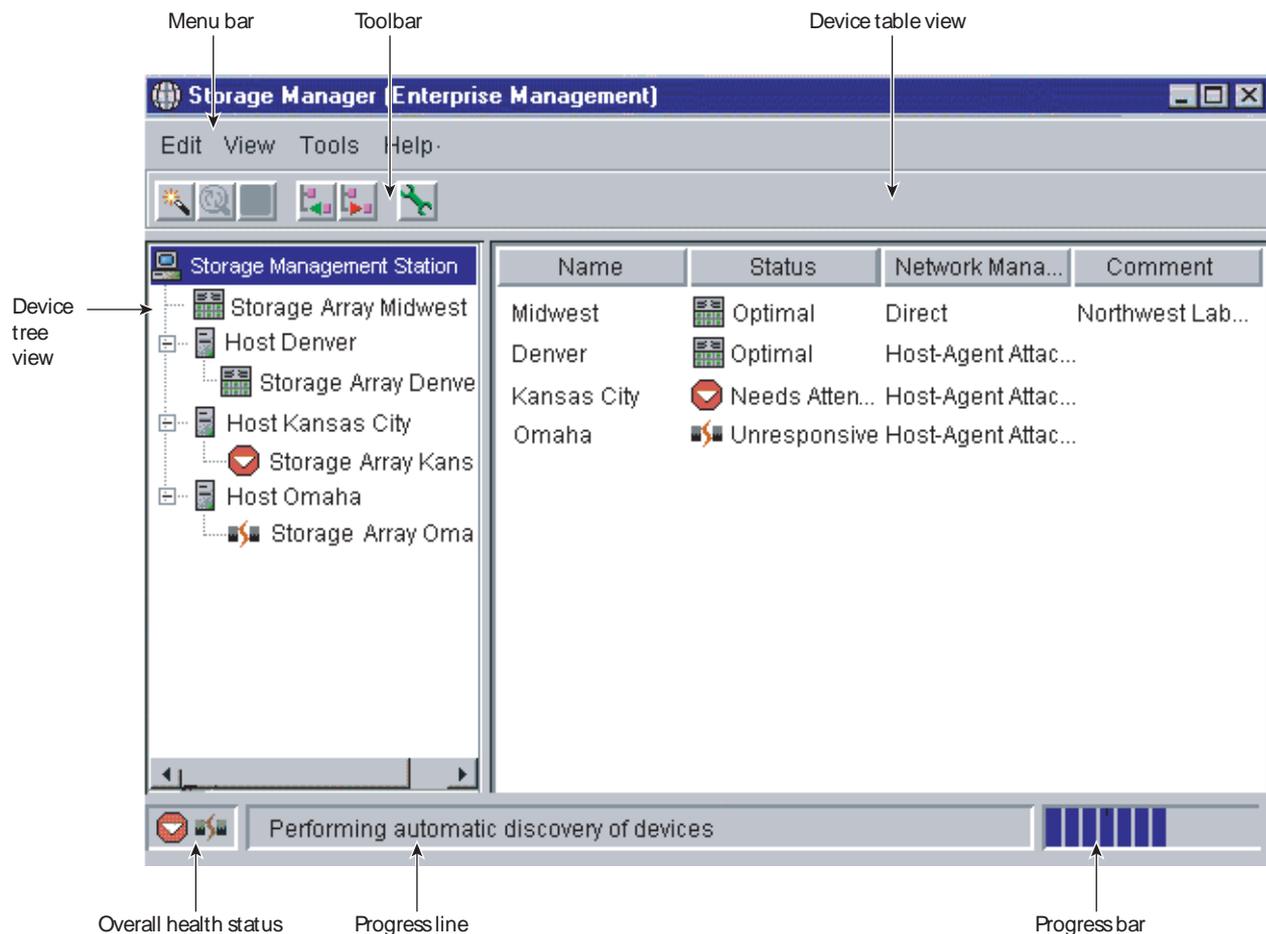


Figure 1-5 Enterprise Management Window

Device Tree

The Device Tree provides a tree-structured view of nodes in the management domain. The storage management station node is the root node and represents the storage management station sending the storage management commands. The Device Tree displays the configured management domain. You add hosts running host-agent software and storage arrays to the management domain, and they are shown in the Device Tree as child nodes of the storage management station node.

The **Enterprise Management** window Device Tree displays the management domain. Figure 1-6 shows an example Device Tree.

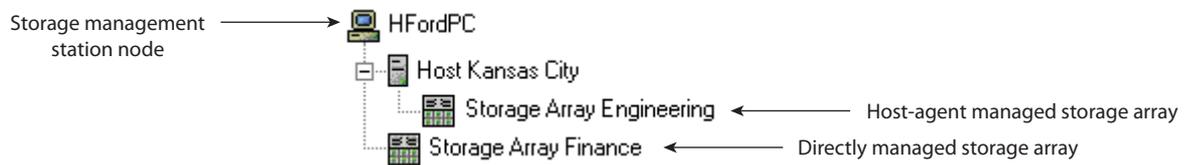


Figure 1-6 Device Tree with a Management Domain

You can use either of two ways to populate a management domain:

1. Use the **Automatic Discovery** option to detect directly managed and host-agent managed storage arrays on the local subnetwork and add them to the management domain. The Enterprise Management window detects host-agent managed storage arrays by first locating the hosts that provide network management connections to the storage arrays. Then the host appears in the Device Tree along with its associated storage arrays.
2. Use the **Add Device** option to manually add directly managed and host-agent managed storage arrays to the management domain. For a directly managed storage array, enter a host name or IP address for each controller in the storage array. For a host-agent managed storage array, enter a name or IP address for the host that is attached to the storage array.

Note: When you discover or add storage arrays to a management domain, they are initially shown as <unnamed> in the Device Tree and Device Table. Select each unnamed storage array, open the **Array Management** window, and use the **Storage Array > Rename** option to name the storage array.

A check appears next to each node in the Device Tree View where alert notification has been set/enabled.

For more information on populating a management domain, see the **Enterprise Management** window Help.

Device Table

The Device Table lists the name, status, management type (direct or host-agent), and comments you have entered for storage arrays. (See Figure 1-5 on page 12.)

For more information about using the Device Table, see the **Enterprise Management** window Help.

Enterprise Management Window Menus

The **Enterprise Management** window menus on the menu bar are described in Table 1-5.

Table 1-5 Enterprise Management Window Menus

Menu	Description
Edit	<p>Allows you to add or remove devices from the management domain, configure alert destinations, or add a comment to the Device Table about a storage array.</p> <p>Important: The storage management software sends critical alerts (using SNMP trap messages) to an NMS destination that you specify using the Edit > Alert > Destinations option in the Enterprise Management window. For the critical alerts to be sent, the Enterprise Management software must be running (the window can be minimized).</p>
View	<p>Allows you to sort the entries in the Device Table by Name, Status, Management Type, or Comment. You can also open the Task Assistant Launch Pad (TALP) from this menu.</p> <p>This menu also shows <i>Partially Managed Devices</i>, which are storage arrays that contain two controllers. However, only one controller of the pair is communicated with to perform storage management operations.</p> <p>For more information on partially managed storage arrays, see the Enterprise Management window Help.</p>
Tools	<p>Allows you to automatically discover devices on the same subnetwork or rescan to find storage arrays newly attached to a host. You can also update the event monitor, open an Array Management window to manage a selected storage array, open the Script Editor to perform batch management tasks, and load a saved configuration file.</p>
Help	<p>Allows you to view the Help system and software version and copyright information for the Enterprise Management software.</p>

For more information about using the Enterprise Management window menu options, see the Enterprise Management window Help.

Enterprise Management Window Toolbar

The **Enterprise Management** window toolbar buttons are described in Table 1-6.

Table 1-6 Enterprise Management Window Toolbar Buttons

Toolbar Button	Description
Automatically detect new devices 	Activates the Automatic Discovery option that detects hosts and storage arrays on the local subnetwork and adds them to the management domain.
Rescan selected host for new devices 	Activates the Rescan option that discovers storage arrays newly attached to the selected host. After physically adding new storage arrays, restart the host-agent software so that it can recognize the new devices before selecting this button. Note: This option is available only when you select a host in the Device Tree.
Display Task Assistant Launch Pad 	Select this toolbar button to display the Task Assistant Launch Pad (TALP). The TALP provides a central dialog from which you can access common administrative tasks, such as adding a storage array or naming a storage array.
Synchronize Event Monitor 	Synchronizes the event monitor with any changes made in the Enterprise Management window, such as adding or removing devices or adding alerts. Note: This option is available only if the configurations of the Enterprise Management window and the event monitor are not synchronized.
Add host/device 	Opens the Add Device dialog for manually adding hosts or storage array controllers to the management domain.

Table 1-6 Enterprise Management Window Toolbar Buttons (**continued**)

Toolbar Button	Description
Remove host/device 	Removes the selected storage array or the selected host and its attached storage arrays from the management domain. Note: This option is available only when you select a storage array in the Device Tree or Device Table, or a host in the Device Tree.
Launch an Array Management window 	Opens an Array Management window for the selected storage array. Note: This option is available only when you select a storage array in the Device Tree or Device Table.

See Figure 1-5 on page 12 for an example of the **Enterprise Management** window toolbar. For more information on using the toolbar options, see the **Enterprise Management** window Help.

Array Management Window

The **Array Management** window is used to configure and maintain the logical and physical components of a storage array and to view and define volume-to-LUN mappings. See Figure 1-7 on page 19.

The **Array Management** window is specific to an individual storage array; therefore, you can only manage a single storage array within an **Array Management** window. However, you can start other **Array Management** windows from the **Enterprise Management** window to manage multiple storage arrays simultaneously.

Note: Depending upon the version of the firmware installed on the storage array, the management software will not be able to manage the storage array until the firmware has been updated. Please review the release notes on the CD-ROM for information about what are the minimum firmware levels requirements.

Note: Depending on your version of storage array software management and firmware, the views, menu options, and functionality may differ from the information presented in this document.

The features of a particular release of firmware will be accessible when an **Array Management** window is launched from the **Enterprise Management** window to manage a storage array. For example, you manage two storage arrays using this software; one storage array has firmware version 5.40 and the other has firmware version 6.12. When you launch an **Array Management** window for a particular storage array, the correct **Array Management** window version is used. The storage array with firmware version 5.40 will use **Array Management** version 08.40.Gx.xx of the storage management software while the storage array with firmware version 6.12 will use **Array Management** version 09.12.Gx.xx. You can verify the version you are currently using by selecting **Help > About** in the **Array Management** window.

This bundling of previous versions of the **Array Management** window provides the flexibility of upgrading the firmware only on selected storage arrays instead of having to perform an upgrade on all storage arrays at once.

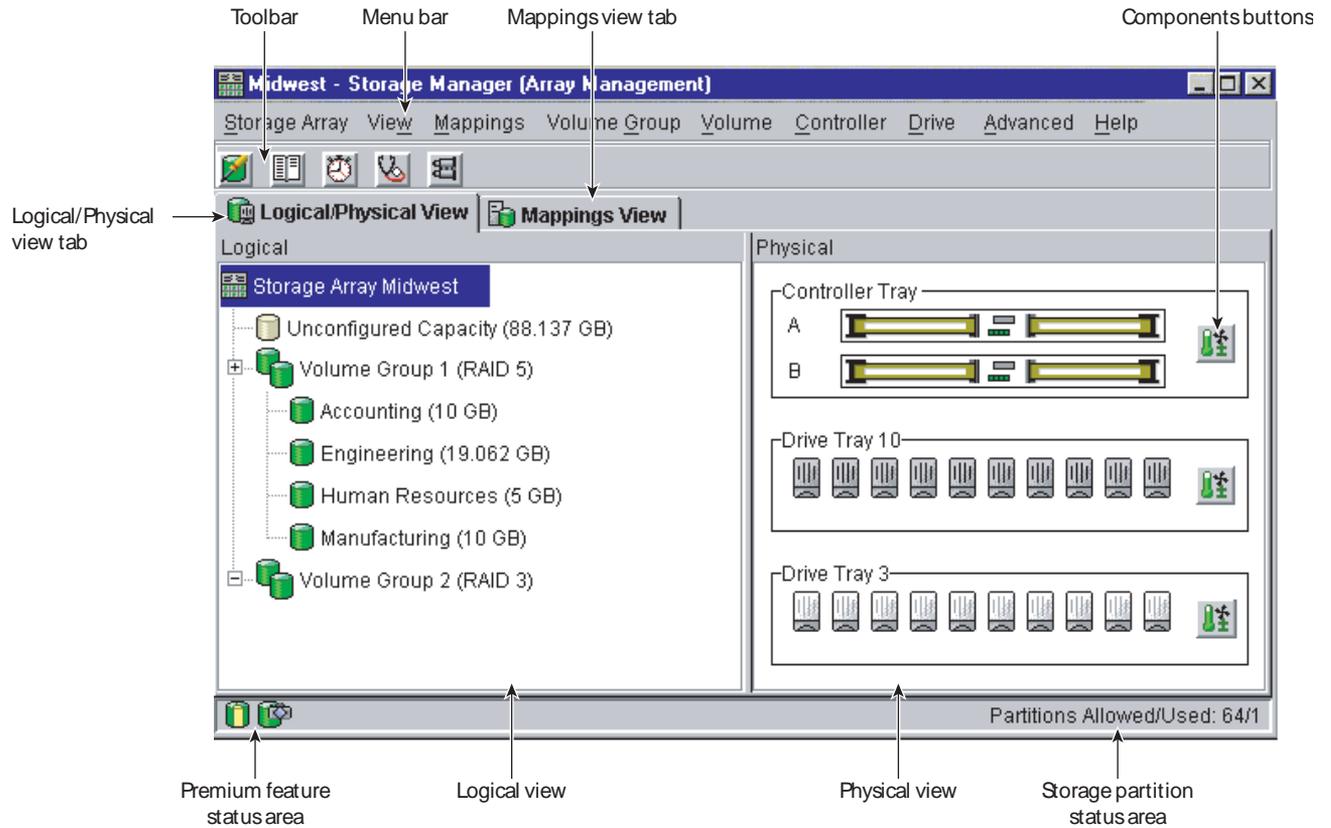


Figure 1-7 Array Management Window

Array Management Window Tabs

The **Array Management** window has two tabs: **Logical/Physical View** and **Mappings View**, which are described in Table 1-7. See Figure 1-7 on page 19. The tabs display information about the logical components (volume and volume groups), physical components (controllers and drives), topological elements (host groups, hosts, host ports), and volume-to-LUN mappings in a storage array.

Table 1-7 Array Management Window Tabs

Tab	Description
Logical/Physical View	<p>The Array Management window Logical/Physical View contains two panes: the Logical view and the Physical view.</p> <p>The Logical view in the left pane of Figure 1-7 on page 19 shows the organization of storage array capacity into volume groups and volumes.</p> <p>The Physical view in the right pane of Figure 1-7 on page 19 shows the physical devices in the storage array, such as command enclosure and drive enclosure components. Selection of a volume or other entity in the Logical view shows you the associated physical components in the Physical view.</p> <p>A Components button in every enclosure presents the status of all enclosure components and shows the temperature status of the enclosure.</p>
Mappings View	<p>The Mappings View of the Array Management window contains two panes: the Topology view and the Defined Mappings view. For details, see “Mappings View” on page 43.</p> <p>The Topology view provides a tree-structured view of logical nodes related to storage partitions.</p> <p>The Defined Mappings detail view displays the mappings associated with the selected node in the Topology view.</p>

For more specific information about the **Logical/Physical View**, **Components** button, or **Mappings View**, see the **Array Management** window Help.

Array Management Window Menu Options

The **Array Management** window menus on the menu bar are described in Table 1-8 on page 21. The menus are used to perform storage management operations for a selected storage array or for selected components within a storage array. See Figure 1-7 on page 19 for an example of the **Array Management** window menus.

Table 1-8 Array Management Window Menus

Menu	Description
Storage Array	Allows you to perform the following storage array management operations: locating functions (locating the storage array or a specific drive channel by flashing indicator lights), configuring the storage array, enabling premium features, starting Recovery Guru, monitoring performance, downloading firmware and NVSRAM files, changing various settings, setting controller clocks, redistributing volumes, running Read Link Status diagnostics, activating or deactivating the remote volume mirroring premium feature, and renaming storage arrays.
View	Allows you to open the Task Assistant Launch Pad (TALP), change the display from the Logical/Physical View to the Mappings View , view Associated Components for a selected node, find a particular node in the Logical or Topology view, locate an associated node in the tree, or access the Event Log or the Storage Array Profile.
Mappings	Allows you to make changes to or retrieve details about mappings associated with a selected node. The options are the following: Define, Change, Move, Replace Host Port, Show All Host Port Information, Delete, and Rename.
Volume Group	Allows you to perform the following volume group storage management operations: creating or locating volumes, changing the RAID level, changing controller ownership, adding free capacity (drives), defragmenting free capacity, placing controllers online or offline, initializing a volume group, reviving a volume group, checking redundancy, or deleting a volume group.
Volume	Allows you to perform the following volume storage management operations: creating volumes; changing ownership or segment size; increasing capacity; initializing, deleting, or renaming a volume; creating or disabling a snapcopy volume; creating a remote volume mirror; or viewing volume properties.
Controller	Allows you to perform the following controller storage management operations: placing a controller online or offline, enabling or disabling data transfer (I/O), changing the controller mode to active, changing the preferred loop ID, running diagnostics, or viewing controller properties.

Table 1-8 Array Management Window Menus (**continued**)

Menu	Description
Drive	Allows you to perform the following drive storage management operations: locating a drive; assigning or unassigning a hot spare; failing, reconstructing, reviving or initializing a drive; or viewing drive properties.
Advanced	Presents maintenance options, which should only be used under the guidance of technical support: Downloading Drive Firmware and Mode Pages, and ESM firmware Capture State Information Reset Controller Persistent Reservations (Not used by SGI)
Help	Displays the Help system, a reference of all Recovery Guru procedures, and software version and copyright information for the Array Management software.

For more information about using the **Array Management** window menu options, see the **Array Management** window Help.

Array Management Window Toolbar Options

The **Array Management** window toolbar buttons are used to create new volumes or volume groups, monitor performance, view events, recover from failures, and locate a

particular node. Each toolbar button is described in Table 1-9. See Figure 1-7 on page 19 for an example of the **Array Management** window toolbar.

Table 1-9 Array Management Window Toolbar Buttons

Toolbar button	Description
Create new volumes (and volume groups, if applicable) 	Allows you to create volume groups and volumes. Note: You must select either a Free Capacity node or an Unconfigured Capacity node in the Logical view before this option is available. If you select a Free Capacity node, the option is called Create New Volumes. If you select an Unconfigured Capacity node, the option is called Create New Volume Groups/Volumes.
View diagnostic event log 	Opens the Event Log Viewer. Note: This log is for customer support use.
Monitor performance 	Opens the Performance Monitor.
Display Task Assistant Launch Pad 	Select this toolbar button to display the Task Assistant Launch Pad (TALP). The TALP provides a central dialog from which you can access common administrative tasks, such as configuring a storage array or defining a host.
Recover from failures 	Opens the Recovery Guru . Note: If the storage array is in the Needs Attention state, the Recovery Guru toolbar button flashes. 
Find node in tree 	Starts the Find dialog that allows you to search for a particular node in the Logical/Physical View or Mappings View of the Array Management window.

For more information on using the toolbar options, see the **Array Management** window Help.

Configuring Storage Arrays

This chapter describes various storage array configuration options to maximize data availability. It also outlines the high-level steps to configure available storage array capacity into volumes and storage partitions.

Volumes and Drive Groups

When you configure a storage system for the first time, you must consider which data protection strategy is most appropriate for your storage system, together with how the total storage capacity must be organized into volumes and shared among hosts.

The storage management software identifies several distinct volumes:

- Standard volumes
- Snapshot volumes
- Snapshot repository volumes
- Primary volumes
- Secondary volumes
- Mirror repository volumes
- Source volumes
- Target volumes

Standard Volumes

Important: The host operating system might have specific limits about how many volumes that the host can access. You must consider these limits when you create volumes that are used by a particular host. For operating system restrictions, refer to the Storage System Product Release Notes that were shipped with the storage management software, and refer to the documentation for your host operating system.

A standard volume is a logical structure that is created on a storage system for data storage. A standard volume that is defined over a set of drives is called a drive group, which has a defined RAID level and capacity. You can create a volume from unconfigured capacity, unassigned drives, or free capacity nodes on the storage system. If you have not configured any volumes on the storage system, the only node available is the unconfigured capacity node.

Use the Create Volume Wizard to create one or more volumes on the storage system. During the volume creation process, the wizard prompts you to select the capacity to allocate for the volumes and to define basic volume parameters and optional advanced volume parameters for the volume.

Drive Groups

A drive group is a set of drives that the controller logically groups together to provide one or more volumes to an application host.

To create a drive group, you must specify two key parameters: the RAID level and the capacity (how large you want the drive group to be). For the capacity parameter, you can either select the automatic choices provided by the software or select the manual method to indicate the specific drives to include in the drive group. Whenever possible, use the automatic method because the storage management software provides the best selections for drive groupings.

Drive Group Creation

The Create Volume Wizard guides you through the steps to create one or more drive groups in a storage system and to configure basic volume parameters and optional volume parameters.

Important: The storage management software determines the default initial capacity selections based on whether you select free capacity, unconfigured capacity, or unassigned drives in the Create Volume Wizard. After the wizard begins, you can change the capacity by defining a new volume capacity in the Specify Capacity/Name dialog.

You can organize available capacity on a storage system by using these types of storage space:

- **Free capacity** – Free capacity is unassigned space in a drive group that you can use to make a volume. When you create a volume from free capacity, an additional volume is created on an existing drive group.
- **Unconfigured capacity** – Unconfigured capacity is available space on drives of a storage system that have not been assigned to a drive group.
- **Unassigned drive** – An unassigned drive is a drive that is not being used in a drive group or is not assigned as a hot spare.

Specifying Volume Parameters

Table 2-1 shows the volume parameters for free capacity, unconfigured capacity, and unassigned drives.

Table 2-1 Volume Parameters for Free Capacity, Unconfigured Capacity, and Unassigned Drives

Parameter	Free Capacity	Unconfigured Capacity	Unassigned Drive
Drive Group Creation	The drive group is predefined.	You must create a drive group before configuring a new volume.	You must create a drive group before configuring a new volume.
Specify Capacity/Name Dialog	Assign a name to the volume. Change the default capacity.	Assign a name to the volume. Change the default capacity.	Assign a name to the volume. Change the default capacity.
SANshare Storage Partitioning will be used	Not applicable	Specify the Map Later option by using the Mappings View. This option specifies that a LUN not be assigned to the volume during volume creation. This option defines specific volume-to-LUN mapping and creates storage partitions.	Specify the Map Later option by using the Mappings View. This option specifies that a LUN not be assigned to the volume during volume creation. This option defines specific volume-to-LUN mapping and creates storage partitions.
SANshare Storage Partitioning will not be used	Not applicable	Use the Default volume-to- LUN option. This option automatically assigns the next available LUN in the default group to the volume. The option grants volume access to host groups or hosts that have no specific volume-to-LUN mappings, which are shown under the Default Group node in the Topology View.	Use the Default volume-to- LUN option. This option automatically assigns the next available LUN in the default group to the volume. The option grants volume access to host groups or hosts that have no specific volume-to-LUN mappings, which are shown under the Default Group node in the Topology View.
Advanced Volume Parameters	You can customize these advanced volume parameters: <ul style="list-style-type: none"> • Custom dynamic cache read prefetch • Segment size • Preferred controller owner • Volume-to-LUN mapping 	You can customize these advanced volume parameters: <ul style="list-style-type: none"> • Custom dynamic cache read prefetch • Segment size • Preferred controller owner • Volume-to-LUN mapping 	You can customize these advanced volume parameters: <ul style="list-style-type: none"> • Custom dynamic cache read prefetch • Segment size • Preferred controller owner • Volume-to-LUN mapping

Protecting Your Data

The storage arrays are designed to be reliable and feature redundant drives, controllers, power supplies, and fans. This hardware redundancy keeps the storage array working if a component fails. In addition, you can use the storage management software to implement the data protection strategies described in the following sections:

- “Software Redundancy” on page 29
- “RAID Levels and Data Redundancy” on page 30
- “Hardware Redundancy” on page 34
- “Storage Array Password Protection” on page 36

For conceptual information and detailed procedures for the strategies describes, see “Learn About Data Protection Strategies” in the Array Management window online help.

Software Redundancy

The storage management software has two premium features that provide data protection strategies. Remote Volume Mirroring is used to create an online, real-time replication of data between storage arrays over a remote distance, while snapshot Volume Copy provides a way to more efficiently copy data either for backup or for application testing.

The following sections provide a brief overview of the premium features used for data protection strategies.

Use Remote Volume Mirroring

When you create a remote volume mirror, a mirrored volume pair is created, which consists of a primary volume on a primary storage array and a second volume on a secondary storage array. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the secondary volume. The secondary volume maintains a mirror (or copy) of the data from its associated primary volume. In the event of a disaster or catastrophic failure at the primary site, a manual role reversal can be performed by the system administrator to promote the secondary volume to a primary role. Hosts will then be able to access the newly promoted volume and business operations can continue.

Create a Snapshot Volume

A snapshot volume is a point-in-time image of a volume. It is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space.

Typically, a snapshot volume is created so that an application, such as a backup application, can access the snapshot volume and read the data while the base volume remains online and user-accessible. When the backup completes, the snapshot volume is no longer needed. You can also create several snapshot volumes of a base volume and write data to the snapshot volumes to perform testing and analysis. For example, before upgrading a database management system, snapshot volumes can be used to test different configurations. The performance data provided by the storage management software can also be used to help decide how to configure the live database system.

RAID Levels and Data Redundancy

RAID is an acronym for Redundant Array of Independent Disks. It describes a storage solution in which the same data or information about the data (parity) is stored in different places on multiple hard disks. By placing data on multiple disks, I/O operations overlap and performance improves. If a disk drive in a volume group fails, the redundant or parity data can be used to regenerate the user data on replacement disk drives.

RAID relies on a series of configurations, called levels, to determine how user and redundancy data is written and retrieved from the drives. This storage management software offers four formal RAID level configurations: RAID levels 0, 1, 3, 5, 6, and 10. Table 2-2 on page 32 describes these configurations. Each level provides different performance and protection features.

RAID levels 1, 3, 5, 6, and 10 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data (mirrored) or an error-correcting code derived from the data. If a drive fails, the redundancy data can be used to quickly reconstruct information.

You configure a single RAID level across a single volume group. All redundancy data for that volume group is stored within the volume group. The capacity of the volume group is the aggregate capacity of the member drives, minus the capacity reserved for redundancy data. The amount of capacity needed for redundancy depends on the RAID level used.

Dynamic RAID-level Migration

Dynamic RAID-Level Migration (DRM) is a modification operation that lets you change the RAID level on a selected drive group. The drive group must contain sufficient free space and a permissible number of drives or the DRM request is rejected.

During the migration process, you can continue to access data on drive groups, volumes, and drives in ISSM-EE. However, you cannot cancel the DRM operation after the process begins.

Note: If you are using dynamic RAID-level migration to migrate a drive group to RAID level 6, your storage array(s) must be running a version of firmware that has support for RAID 6 before you start the Dynamic RAID-level Migration.

Table 2-2 RAID Level Configurations

RAID Level	Short Description	Detailed Description
RAID 0	Non-redundant, Striping Mode	<p>RAID 0 offers high performance, but does not provide any data redundancy. It stripes data across all of the drives in the volume group.</p> <p>RAID 0 is not recommended for high availability needs. It is better for non-critical data.</p> <p>If a single drive fails in the volume group, all associated volumes fail and all data is lost.</p>
RAID 1 or RAID 10	Striping/Mirroring Mode	<p>RAID 1 offers high performance and uses disk mirroring to make an exact copy from one drive to another drive.</p> <p>When more than 2 drives are used, RAID 1 is also called RAID 10 or 0+1 (indicating striping and mirroring). In this situation, data is striped across half the drives, which are then mirrored. For example, when there are 4 drives, data is striped across two drives and mirrored on the other two. Because the data is mirrored, you can never have an odd number of disks in a RAID member.</p> <p>RAID 1 offers the best data availability, but only half of the drives in the volume group are available for user data.</p> <p>If a single drive fails in a RAID 1 volume group, all associated volumes become degraded, but the mirror drive still allows the data to be accessed.</p> <p>RAID 1 can survive multiple drive failures as long as no more than one failure exists per mirrored pair. If a drive-pair fails in a RAID 1 volume group, all associated volumes fail and all data is lost.</p>
RAID 3	High Bandwidth Mode	<p>User data and redundant information (parity) is striped across the drives. The equivalent of one drive's worth of capacity is used for redundant information.</p> <p>RAID 3 is good for large data transfers in applications such as multimedia or medical imaging that write and read large sequential chunks of data.</p> <p>If a single drive fails in a RAID 3 volume group, all associated volumes become degraded, but the redundant information still allows the data to be accessed. If two or more drives fail in a RAID 3 volume group, all associated volumes fail and all data is lost.</p>

Table 2-2 RAID Level Configurations (**continued**)

RAID Level	Short Description	Detailed Description
RAID 5	High I/O Mode	<p>User data and redundant information (parity) is striped across the drives. The equivalent of one drive's worth of capacity is used for redundant information.</p> <p>RAID 5 is good for multi-user environments such as database or filesystem storage in which typical I/O size is small and there is a high proportion of read activity.</p> <p>If a single drive fails in a RAID 5 volume group, all associated volumes become degraded, but the redundant information still allows the data to be accessed. If two or more drives fail in a RAID 5 volume group, all associated volumes fail and all data is lost.</p>
RAID 6	High I/O mode with simultaneous drive failure protection	<p>RAID 6 stripes both user data and redundancy data (parity) across the drives and uses the equivalent of the capacity of two drives (in a drive group) for redundancy data. It protects against the simultaneous failure of two drive group member drives by using two independent error-correction schemes. RAID 6 is not supported on the TP9700/IS4500 storage systems.</p>

Hardware Redundancy

Data protection strategies provided by the storage system hardware include cache memory, hot spare drives, background media scans, and channel protection.

Protect Data in the Controller Cache Memory

Caution: Possible loss of data access – Do not enable the “Write caching without batteries” parameter unless you have a uninterruptible power supply (UPS) for power loss protection. If you have the parameter enabled and your batteries are low or discharged, the data in the cache could be lost if a power outage occurs.

Write caching can increase I/O performance during data transfers. Caching also increases the risk for data loss if a controller (or its memory) fails while unwritten data resides in cache memory. Use write cache mirroring to protect data during a controller or cache memory failure. When you enable *write cache mirroring*, cached data is mirrored across two redundant controllers with the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. Enable the Write Cache Mirroring parameter for each volume by using the **Volume > Properties** option in the **Array Management** window.

To prevent data loss or corruption, the controller periodically writes cache data to disk (flushes the cache). When the amount of unwritten data in cache reaches a certain level, called a *start percentage*, or when data has been in cache for a predetermined amount of time, the software signals the controller to write the data to disk. The controller writes to disk until the amount of data in cache drops to a *stop percentage* level. Start and stop percentages can be configured by the user. For example, you can specify that the controller start flushing the cache when the cache reaches 80% full and stop flushing the cache when the cache reaches 16% full.

For maximum data protection, you can choose low start and stop percentages. However, in both cases this increases the chance that data needed for a host read will not be in the cache, which decreases the cache hit percentage (described in Table 2-7 on page 70), and therefore the I/O request rate (also described in Table 2-7 on page 70). Choosing low start and stop percentages also increases the number of disk writes necessary to maintain the cache level, increasing system overhead and further decreasing performance.

If a power outage occurs, data in the cache that has not been written to the drive media will be lost, even if it is mirrored to the cache memory of both controllers. For protection against this occurrence, batteries in the controller enclosures and array enclosures protect the cache against power outages. Change the controller enclosure and array enclosure batteries at the recommended time intervals. The ISSM-EE storage management software features a battery age clock you can set when you replace a battery. This clock keeps track of the age (in days) of the battery so you know when it is time to replace the battery. You will receive a critical alert notification when the battery is nearing expiration and when it has reached expiration.

Configure Hot Spare Drives

A valuable strategy to keep data available is to assign available drives in the storage array as hot spare drives. A *hot spare* is a drive containing no data that acts as a standby in the storage array in case a drive fails in a RAID 1, 3, 5, 6 or 10 volume. The hot spare drive adds another level of protection to the storage array.

If a drive fails in the storage array, the hot spare is automatically substituted for the failed drive without requiring a physical swap. If the hot spare drive is available when a drive fails, the controller will use redundancy data to reconstruct the data from the failed drive to the hot spare drive. When you have physically replaced the failed drive, you can use either of the following options to restore the data:

- When you have replaced the failed disk drive, the data from the hot spare is copied back to the replacement disk drive. This is called *copyback*.
- You can assign the hot spare as a permanent member of the volume group. Performing the copyback function is not required for this option.

If you do not have a hot spare, you can still replace a failed disk drive while the storage array is operating. If the drive is part of a RAID 1, 3, 5, 6, or 10 volume group, the controller will use redundancy data to automatically reconstruct the data from the failed drive to the replacement drive. This is called *reconstruction*.

Channel Protection

In a Fibre Channel environment, channel protection is usually present for any storage system. When the storage system is cabled correctly, two redundant arbitrated loops (ALs) exist for each drive.

Enable a Background Media Scan

A *media scan* is a background process performed by the controllers to provide error detection on the drive media. A media scan detects errors and reports them to the Event Log.

The media scan must be enabled for the entire storage array as well as enabled on each volume.

The media scan runs on all volumes in the storage array for which it has been enabled. The advantage of enabling the media scan process is that the process can find media errors before they disrupt normal drive reads and writes. The media scan process scans all volume data to verify that it can be accessed and, if you enable a *redundancy check*, scans the volume redundancy data at the same time.

Storage Array Password Protection

Caution: Running operations that alter the configuration of your storage system can cause serious damage, including data loss. Configuring a password for each storage system that you manage prevents unauthorized access to destructive commands.

The storage management software provides other security features to protect data, including generation numbering to prevent replay attacks and hashing and encryption to guard against client spoofing and snooping.

For added security, you can configure each storage system with a password to protect it from unauthorized access. A password protects any options that the controller firmware deems destructive. These options include any functions that change the state of the storage system, such as creating a volume or modifying the cache setting.

Note: If you forget the password, contact technical support.

After the password has been set on the storage array, you will be prompted for that password the first time you attempt a destructive operation in the Array Management Window. You will be asked for the password only once during a single management session.

For storage systems with a password and alert notifications configured, any attempts to access the storage system without the correct password are reported.

Password Failure Reporting and Lockout

For storage arrays with password and alert notifications configured, any attempts to access the storage array without the correct password will be reported. If a password is entered incorrectly, an information major event log (MEL) event is logged, indicating that an invalid password or no password has been entered.

If the password is entered incorrectly 10 times within 10 minutes, both controllers will enter lockout mode. The lockout mode will last for 10 minutes, during which time both controllers will deny any attempts to enter a password to access the storage array.

Note: if the controllers are reset, the password failure counter will be cleared and access to the storage array can be attempted again. If the password is entered incorrectly after 10 minutes, the controllers will re-enter lockout mode.

A critical MEL event will be logged to the event log, indicating that the controllers have entered lockout mode. After the 10 minute lockout period has elapsed, the controllers will reset the password failure counter and unlock themselves.

I/O Data Path Protection

Input/output (I/O) data path protection to a redundant controller in a storage array is accomplished with the Auto-Volume Transfer (AVT) feature and the multi-path I/O (MPIO) driver in the Windows operating system (OS) and Redundant Dual Active Controller (RDAC) multi-path driver in the Linux OS.

AVT is a built-in feature of the controller firmware that permits ownership of a volume to be transferred to a second controller if the preferred controller fails. When you use AVT with a multi-path driver, AVT helps to make sure that an I/O data path always is available for the volumes in the storage array.

If a component such as a controller, a cable, or an environmental services monitor (ESM) fails, or an error occurs on the data path to the preferred controller, AVT and the

multi-path driver automatically transfer the volume groups and volumes to the alternate “nonpreferred” controller for processing. This failure or error is called a failover.

Multi-path drivers such as MPIO and RDAC are installed on host computers that access the storage array and provide I/O path failover. The multi-path driver (MPIO in the Windows OS and RDAC in the Linux OS) is used for failover. The AVT feature is used specifically for single-port cluster failover. The AVT feature mode is automatically selected by host type.

During a failover, the volume transfer is logged as a critical event. You can configure the alert destinations for the storage array to send an alert notification automatically.

Multi-Path Driver with AVT Enabled

Enabling AVT on your storage array and using it with a host multi-path driver helps to ensure that an I/O data path is always available for the storage array volumes.

When you create a volume on a storage array where AVT is enabled, a controller must be assigned to own the volume, called the preferred controller. The preferred controller normally receives the I/O requests to the volume. If a problem along the data path, such as a component failure, causes an I/O request to fail, the multi-path driver sends the I/O to the alternate controller.

Note: You should have the multi-path driver installed at all times. You should always enable the ATV mode. Set the ATV mode to a single port cluster host type.

After the I/O data path problem is corrected, the preferred controller automatically reestablishes ownership of the volume as soon as the multi-path driver detects the path is normal again.

Multi-Path Driver with AVT Disabled

When you disable AVT on your storage array, the I/O data path is still protected as long as a multi-path driver is installed on each host connected to the storage array. However, when an I/O request is sent to a specified volume, and a problem occurs along the data path to its preferred controller, **all** volumes on the preferred controller are transferred to the alternate controller, not just the specified volume.

Creating Volumes and Storage Partitions

When configuring a storage array, consider appropriate data protection strategies as well as how the total storage capacity will be organized into volumes and shared among hosts in the storage enterprise.

Volume Groups

A volume group is a set of disk drives that the controller logically groups together to provide one or more volumes to an application host. When you create a volume from free capacity, an additional volume is created on an existing volume group.

To create a volume group, you must specify two key parameters: the RAID level and the capacity (how large you want the volume group). For the capacity parameter, you can either select the automatic choices provided by the software or select the manual method to indicate the specific disk drives to include in the volume group. Whenever possible, use the automatic method because the software provides the best selections for disk drive groupings.

Standard Volumes

Important: Potential loss of data – The host operating system might have specific limits on how many volumes the host can access. You must consider these limits when you create volumes that are used by a particular host. For operating system restrictions, refer to the *Storage System Product Release Notes* that were shipped with the software, and refer to your host operating system documentation.

A standard volume is a logical structure created on a storage array for data storage. A volume that is defined over a set of disk drives is called a volume group and has a defined RAID level and capacity. You can create a volume from either unconfigured capacity or free capacity nodes on the storage array. If you have not configured any volumes on the storage array, the only node available is the unconfigured capacity node.

Use the Create Volume Wizard to create one or more volumes on the storage array. During the volume creation process, the Wizard prompts you to select the capacity to allocate for the volumes and to define basic and optional advanced volume parameters for the volume.

Specifying Volume Parameters from Free Capacity

Important: The storage management software determines the default initial capacity selections based on whether you select free capacity, unconfigured capacity, or unassigned disk drives in the Wizard. After the Wizard begins, you can change the capacity by defining a new volume capacity in the Specify Capacity/Name dialog.

In the Create Volume Wizard, the Specify Capacity/Name dialog enables you to define the capacity of the volume and the volume name. The Create Volume Wizard also gives you the option to use the recommended advanced volume parameters for the new volume or to customize the advanced volume parameters for the new volume. Advanced volume parameters that you can customize include the volume I/O characteristics type and the preferred volume ownership.

After defining the capacity and volume name, you can specify the advanced volume I/O characteristics that apply to the volume. These advanced parameters are based on the needs of your application and include custom dynamic cache read prefetch and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter. By specifying these advanced parameters, you can tailor volumes to meet any specific application needs.

Specifying Volume Parameters from Unconfigured Capacity

Important: The free capacity, unconfigured capacity, or unassigned disk drives selected when starting the Wizard determine the default initial capacity selections. After the Wizard begins, you can change the capacity by defining a new volume capacity in the Specify Capacity/ Name dialog.

The Create Volume Wizard will direct you to create a Volume Group. After the Volume Group is created you can specify the RAID level of the volume group to meet your volume data storage and protection requirements, and select the disk drives that will be in the volume group.

Next, you specify the capacity for the volume and volume name and whether to use recommended advanced volume settings or customize the advanced volume properties for this volume. The Create Volume Wizard also gives you the option to use the recommended advanced volume parameters for the new volume or to customize the advanced volume parameters for the new volume. Advanced volume parameters that you can customize include the volume I/O characteristics type and the preferred volume ownership.

Then you specify the volume I/O characteristics that apply to the volume based on the needs of your application, or enable dynamic cache read ahead prefetch and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter. You can tailor volumes to specific application needs by customizing the advanced volume settings.

During the volume-creation process, you are prompted to set the volume-to-LUN mapping preference to specify whether you will use SANshare Storage Partitioning. There are two settings:

- **Default** – If you are not using SANshare Storage Partitioning, specify this setting. The Default setting specifies that a LUN be automatically assigned to the volume using the next available LUN within the default group. This setting grants volume access to host groups or hosts that have no specific volume-to-LUN mappings, which are displayed under the Default Group node in the Topology View.
- **Map later using the Mappings View**– If you are using SANshare Storage Partitioning, specify this setting. The Map later setting specifies that a LUN not be assigned to the volume during volume creation. This setting allows definition of a specific volume-to- LUN mapping and creation of storage partitions.

Register the Volume with the Operating System

After you have created all of your volumes and have assigned volume-to-LUN mappings, use the host failover driver utility to scan the mapped volumes, such as the `hot_add` command when using RDAC, or `scsiha` and `ioconfig` utilities to register the volume with the operating system.

You can then run the `hot_add` utility to ensure that the operating system is aware of the newly created volumes. If available for your operating system, you can run the host-based `smeedevices` utility to associate the physical storage array name and the volume name. See “Storage Partitioning” on page 47 for more information.

Configuring Storage Array Capacity into Volumes

You can configure volumes from either unconfigured capacity or free capacity in a storage array as described in Table 2-3.

Table 2-3 Creating Volumes from Unconfigured or Free Capacity

If you create the volume from...	Then perform this activity...
Unconfigured Capacity	<p>First define the volume group, specifying two parameters: RAID level and capacity. (Capacity refers to how large you want the volume group to be, that is, how many drives to include in the volume group.)</p> <p>Next, specify the volume parameters for the first volume on the volume group. Parameters include volume name, expected usage, segment size, desired capacity, controller ownership, and volume-to-LUN mapping preference.</p> <p>The expected usage parameter allows you to select an optimal segment size for the application that will use this storage volume. If you select either file system, database, or multimedia, and then you select Based on Expected Usage for segment size, you can optimize the segment size without knowing which specific segment size value to set.</p>
Free Capacity	<p>Because free capacity resides only on an existing volume group, the overall capacity (that is, the number of drives) and RAID level were already specified.</p> <p>Therefore, you only need to specify the volume parameters (volume name, usage, desired capacity, controller ownership, and volume-to-LUN mapping preference).</p> <p>Note: You can add free capacity to a volume group by using the Volume Group > Add Free Capacity option.</p>

Mappings View

The **Mappings View** tab is used to define the storage topology elements (host groups, hosts, host ports, and so on), to define volume-to-LUN mappings, and to view Storage Partitioning and heterogeneous host information. The **Mappings View** tab has two views, **Topology** view and **Defined Mappings** view, shown in Figure 2-1 and described in Table 2-4.

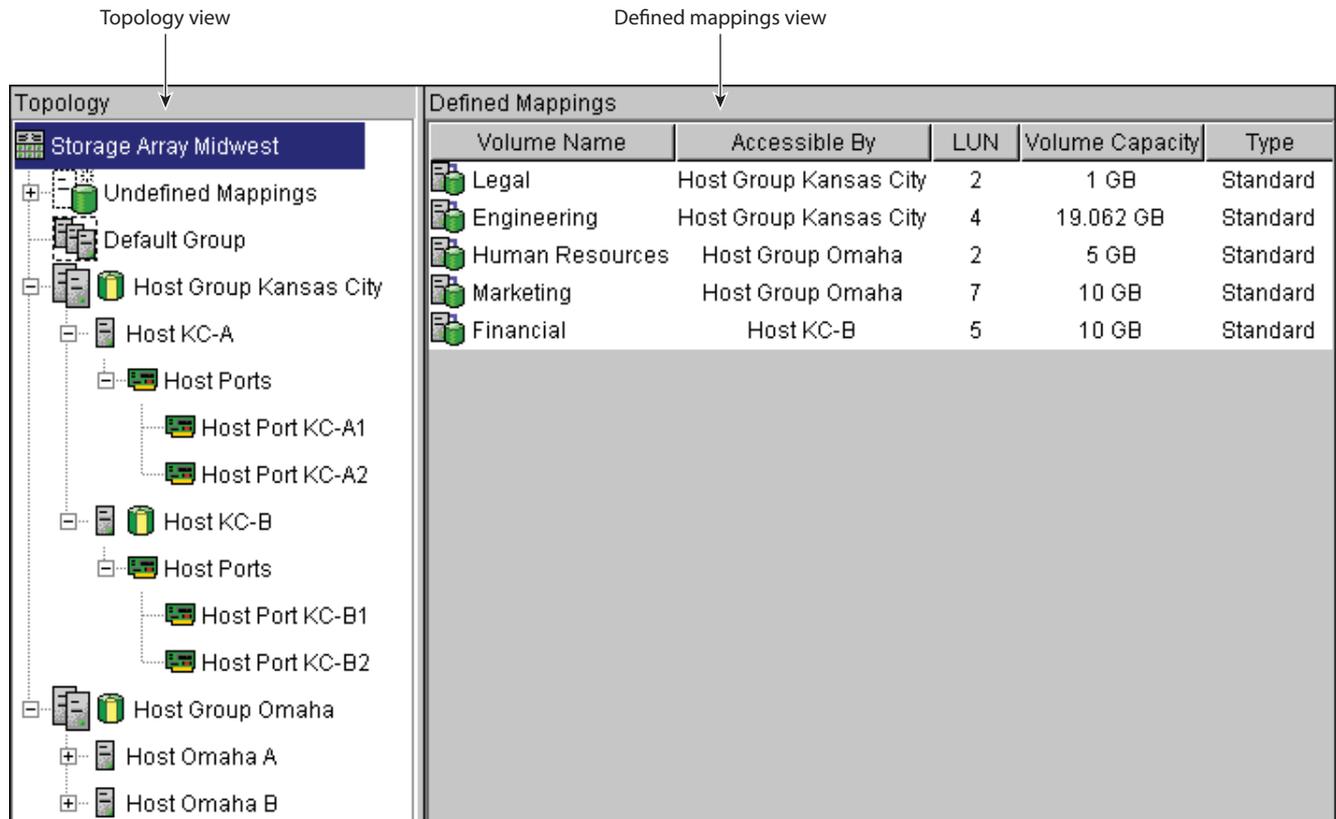


Figure 2-1 Mappings View Window

Table 2-4 Mappings View Tab

View	Description
Topology	Shows defined topological elements (host groups, hosts, and host ports), undefined mappings (volumes that have been created but do not have a defined volume-to-LUN mapping), and the Default Group.
Defined Mappings	Displays the volume-to-LUN mappings in a storage array in table form. Information is displayed about the volumes: topological entities that can access the volume, volume name, volume capacity, and LUN number associated with the volume.

Table 2-5 on page 44 describes the topological elements displayed in the **Mappings View** window in Figure 2-1 on page 43.

Table 2-5 Volume-to-LUN Terminology

Term	Description
Storage Partitioning Topology	A collection of nodes (default group, host groups, hosts, and host ports) shown in the Topology view of the Mappings View tab. You must define the various topological elements if you want to define specific volume-to-LUN mappings and storage partitions for host groups or hosts.
Default Group	A node in the Topology view that designates all host groups, hosts, and host ports that: (1) have no specific volume-to-LUN mappings and (2) share access to any volumes that were automatically assigned default LUN mappings by the controller firmware during volume creation.
Host Group	An optional topological element that you define if you want to designate a collection of hosts that will share access to the same volumes. The host group is a logical entity.
Host	A computer that is attached to the storage array and accesses various volumes on the storage array through its host ports (host bus adapters). You can define specific volume-to-LUN mappings to an individual host or assign the host to a host group that shares access to one or more volumes using the Define Host Wizard.

Table 2-5 Volume-to-LUN Terminology (**continued**)

Term	Description
Host Port	<p>The physical connection that allows a host to gain access to the volumes in the storage array. When the host bus adapter only has one physical connection (host port), the terms host port and host bus adapter are synonymous.</p> <p>Host ports can be automatically detected by the storage management software after the storage array has been connected and powered-up. Therefore, if you want to define specific volume-to-LUN mappings for a particular host or create storage partitions, you must define the host's associated host ports.</p> <p>Initially, all detected host ports belong to the Default Group. Therefore, if during volume creation, you had a LUN automatically assigned to a volume, that volume will be accessible by any of the host ports in the Default Group. If you have the Storage Partitioning feature enabled, then you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation.</p> <p>Use the host bus adapter utility to find the World Wide Name (WWN) of the host port. (This is the host port identifier shown in the Define New Host Port dialog in the Mappings View.) The WWNs for the host ports on a particular host are used to define the host ports and associate them with a particular host using the Define New Host Port dialog. If necessary, see your operating system or host bus adapter documentation for more information.</p> <p>If a host port is moved, any volume-to-LUN mappings must be re-mapped. Access to your data will be lost until this is done.</p>
Logical Unit Number (LUN)	<p>The number a host uses to access a volume on a storage array. Each host has its own LUN address space. Therefore, the same LUN may be used by different hosts to access different volumes on the storage array. However, a volume can only be mapped to a single LUN. A volume cannot be mapped to more than one host group or host.</p> <p>For example, Figure 2-5 on page 54 shows that Host Group KC may access Volume Legal using LUN 2 and Host Group Omaha may access Volume H Resources also using LUN 2.</p>

Table 2-5 Volume-to-LUN Terminology (**continued**)

Term	Description
Default Volume-to-LUN mapping	<p>During volume creation, you can specify that you want to have the software assign a LUN automatically to the volume or that you want to map a LUN to the volume later. If you have the Storage Partitioning feature enabled, then you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation. Any volumes that are given automatic (default) volume-to-LUN mappings can be accessed by all host groups or hosts that do not have specific volume-to-LUN mappings. These host groups and hosts are shown as part of the Default Group in the Topology view of the Mappings View.</p>
Specific Volume-to-LUN mapping	<p>A specific volume-to-LUN mapping occurs when you select a defined host group or host in the Topology view and select the Storage Partitioning Wizard or Define Additional Mapping option to assign a volume a specific LUN (volume-to-LUN mapping). This designates that only the selected host group or host has access to that particular volume through the assigned LUN. You can define one or more specific volume-to-LUN mappings for a host group or host. The Storage Partitioning feature must be enabled to create specific mappings.</p>

Storage Partitioning

Storage Partitioning is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The SANshare Storage Partitioning feature enables hosts with different operating systems (heterogeneous hosts) to share access to a storage array.

A storage partition is a logical entity consisting of one or more storage array volumes that can be shared among hosts. To create a storage partition after the total storage capacity has been configured into volumes, you must define a single host or collection of hosts (or host group) that will access the storage array. Then you must define a volume-to-LUN mapping, which lets you specify the host group or the host that will have access to a particular volume in your storage array. Based on the premium feature key purchased, the storage management software can support a maximum of 64 storage partitions.

You can define a maximum of 255 volumes per partition; this number is limited to the total number of volumes on your storage array. The storage management software can further support up to two host ports in each host and up to eight ports in each host group, allowing a four-way cluster of dual-adapter hosts.

Storage partitions can quickly be created with the **Storage Partitioning Wizard**. The Wizard contains the major steps required to specify which hosts, volumes, and associated logical unit numbers (LUNs) will be included in the partition.

The software can support up to two host ports in each host and up to eight ports in each host group, allowing a four-way cluster of dual-adapter hosts. After the total storage capacity is configured into volumes, those volumes need to be shared among hosts in the storage array. Storage Partitioning allows hosts with the same or different operating systems to share access to a storage array.

Storage Partitioning Example

In the example shown in Figure 2-2, four hosts (Omaha A and B and KC-A and B) are connected to Storage Array Midwest. Host KC-B has exclusive access to Volume Financial. Hosts KC-A and KC-B (Host Group Kansas City) share access to Volumes Legal and Engineering, and Omaha A and Omaha B (Host Group Omaha) share access to Volumes Marketing and H Resources. So, three storage partitions were configured.

The first partition is composed of Volume Financial. This volume is accessed by Host KC-B using LUN 5. Even though Host KC-B is part of the logical Host Group Kansas City,

Host KC-A cannot access this volume because the volume-to-LUN mapping was created with Host KC-B rather than the Host Group Kansas City.

The second partition consists of Volumes Legal and Engineering. This volume-to-LUN mapping was created using Host Group Kansas City. These volumes are accessed by Hosts KC-A and KC-B in Host Group Kansas City using both LUNs 2 and 4.

The third partition consists of Volumes Marketing and H Resources. This volume-to-LUN mapping was created using Host Group Omaha. These volumes are accessed by Hosts Omaha A and Omaha B in Host Group Omaha using both LUNs 7 and 2.

Note: A host gains access to the volumes on the storage array through the physical host ports residing on the installed host bus adapters. To ensure redundant paths to each volume, it is recommended that each host have at least two host ports.

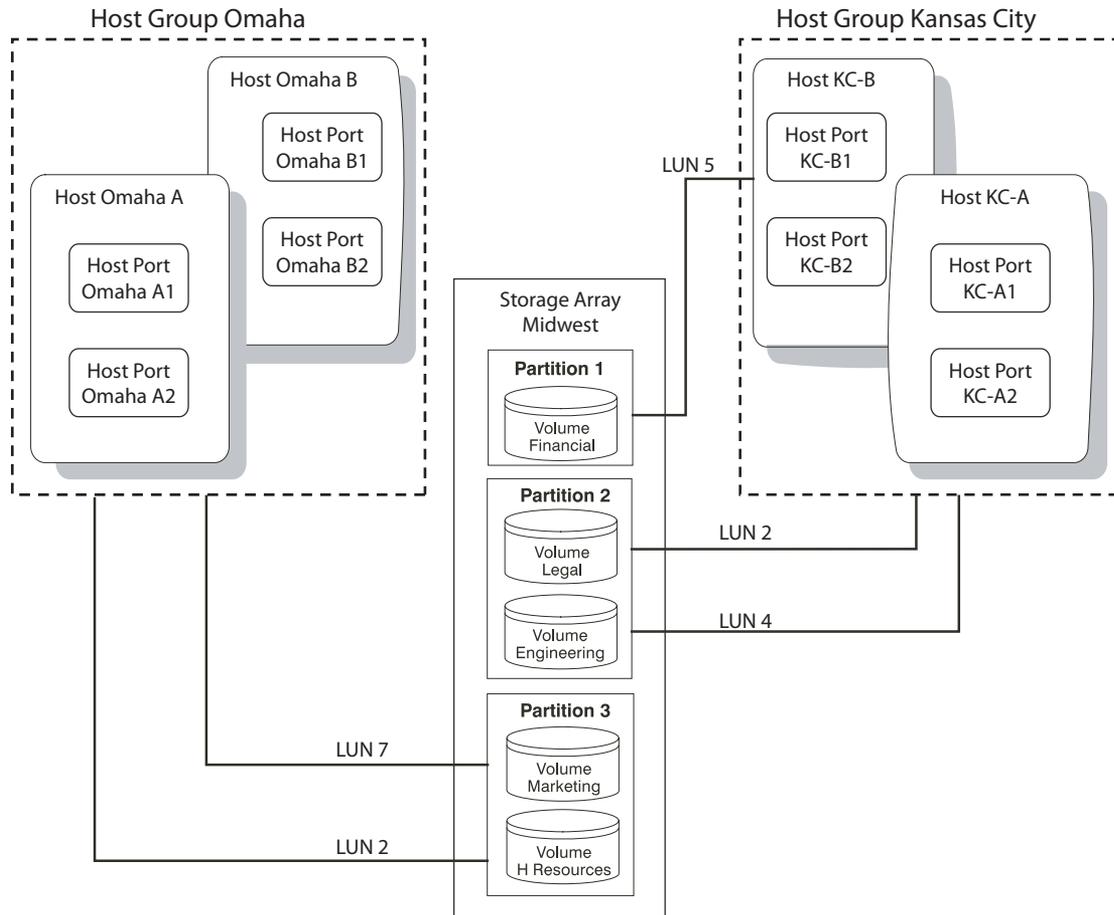


Figure 2-2 Storage Partitioning Example

The next section explains how this example was configured.

Major Steps to Create Storage Partitions

Creating storage partitions involves the following three major steps:

1. Create volumes on the storage array.
2. Define storage partitions topology (host groups, hosts, and host ports).
3. Define volume-to-LUN mappings.

Note: The three major steps provide a basis for understanding what storage partitions are and how to create them. For specific, detailed procedures, see the **Array Management** window Help system.

Create Volumes on the Storage Array

As part of the volume creation, specify one of two volume-to-LUN mapping settings:

- **Automatic** – If you are **not** using Storage Partitioning, specify this setting. The Automatic setting specifies that a LUN be automatically assigned to the volume using the next available LUN within the Default Group. This setting will grant volume access to host groups or hosts that have no specific volume-to-LUN mappings (designated by the Default Group in the Topology view).
- **Map later with Storage Partitioning** – If you **are** using Storage Partitioning, specify this setting. The Map later setting specifies that a LUN not be assigned to the volume during volume creation. This setting allows definition of a specific volume-to-LUN mapping and creation of storage partitions.

Define Storage Partitions Topology

Storage partition topology is reconfigurable. You can do the following:

- Move a host port
- Replace a host port
- Move a host from one host group into another host group
- Delete a host group, host, or host port
- Rename a host group, host, or host port
- Change a volume-to-LUN mapping
- Define additional volume-to-LUN mappings

Define Volume-to-LUN Mappings

The last step is to grant access to one or more volumes to a defined host group or host. You do this by defining specific volume-to-LUN mappings using the **Storage Partitioning Wizard**. Each host group or host is granted a unique view of partitioned storage.

Note: Only one volume-to-LUN mapping is allowed per volume. To grant volume access to more than one host, you must first group the hosts together into a logical host group. Then you can grant volume access to that host group, thereby allowing the hosts in the host group to share access to the volume.

Because each host has its own LUN address space, you can use the same LUN in more than one volume-to-LUN mapping, as long as that LUN is available for use by each particular host participating in the mapping.

You can only define new volume-to-LUN mappings for volumes that currently have no mappings (that is, they do not have default mappings or specific mappings).

Volume-to-LUN mappings are shared between controllers in the storage array.

Each host group or host that you define is granted a unique view of partitioned storage. A defined host group or host can access either:

- Volumes with default volume-to-LUN mappings – The host group or host is part of the Default Group.
- Volumes to which they have been granted access through a specific volume-to-LUN mapping – The host group or host will be part of a storage partition.

In the example in Figure 2-3, none of the topological elements have had specific volume-to-LUN mappings defined. Therefore, they appear under (as part of) the default host group and can access any volumes that have been assigned default LUNs (default volume-to-LUN mappings).

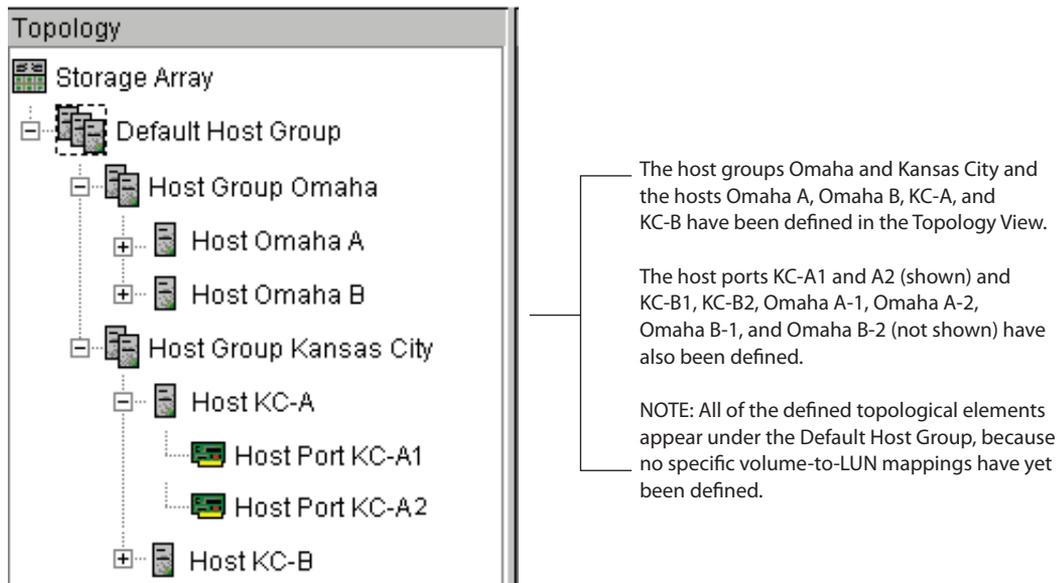


Figure 2-3 Defining Hosts and Host Groups in the Topology View

If you define specific volume-to-LUN mappings for any of the host groups or hosts, they no longer appear under the default host group; they move to under the root node, marked storage array. Also, they can no longer access volumes with default volume-to-LUN mappings. Either a host group or host can access volumes with default mappings or specific mappings, but not both.

Heterogeneous Hosts

The heterogeneous hosts portion of the Storage Partitioning feature allows hosts running different operating systems to access a single storage array. To specify different operating systems for attached hosts, you must specify the appropriate host type when you define the host ports for each host.

Host types can be completely different operating systems, such as Solaris and Windows NT, or variants of the same operating system, such as Windows NT clustered and Windows NT non-clustered. When a host type is specified, it allows the controllers in the

storage array to tailor their behavior (such as LUN reporting and error conditions) depending on the operating system or variant of the host sending the information.

Heterogeneous Hosts Example

Note: Heterogeneous host settings are only available with Storage Partitioning enabled. In Figure 2-5 on page 54, the Storage Partitioning feature is enabled.

In a heterogeneous environment, you must set each host type to the appropriate operating system during host port definition (Figure 2-4). By doing this, the firmware on each controller can respond correctly for that host's operating system.

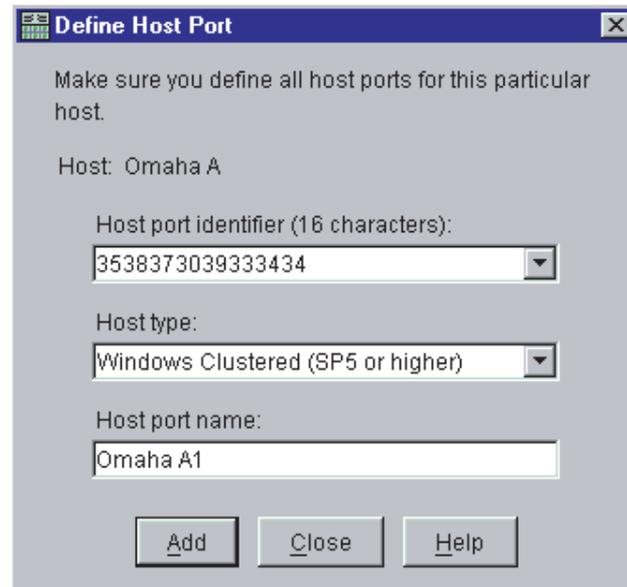


Figure 2-4 Host Port Definitions Dialog

In the example shown in Figure 2-5 on page 54, heterogeneous hosts have been defined on Storage Array Midwest. Four hosts (Omaha A and B and KC-A and B) are connected to Storage Array Midwest. Hosts Omaha A and Omaha B share access to Volumes Marketing and HResources. Host KC-A has exclusive access to Volumes Legal and Engineering, and Host KC-B has exclusive access to Volume Financial. So, three storage

partitions have been configured. Because there are four hosts running three different operating systems, the appropriate host types must be defined for each host port to support these heterogeneous hosts.

After you define the host type for each host port, you can display the host port's host type in the **Topology** view by placing your cursor over the specific host port; a tooltip will display the associated host type.

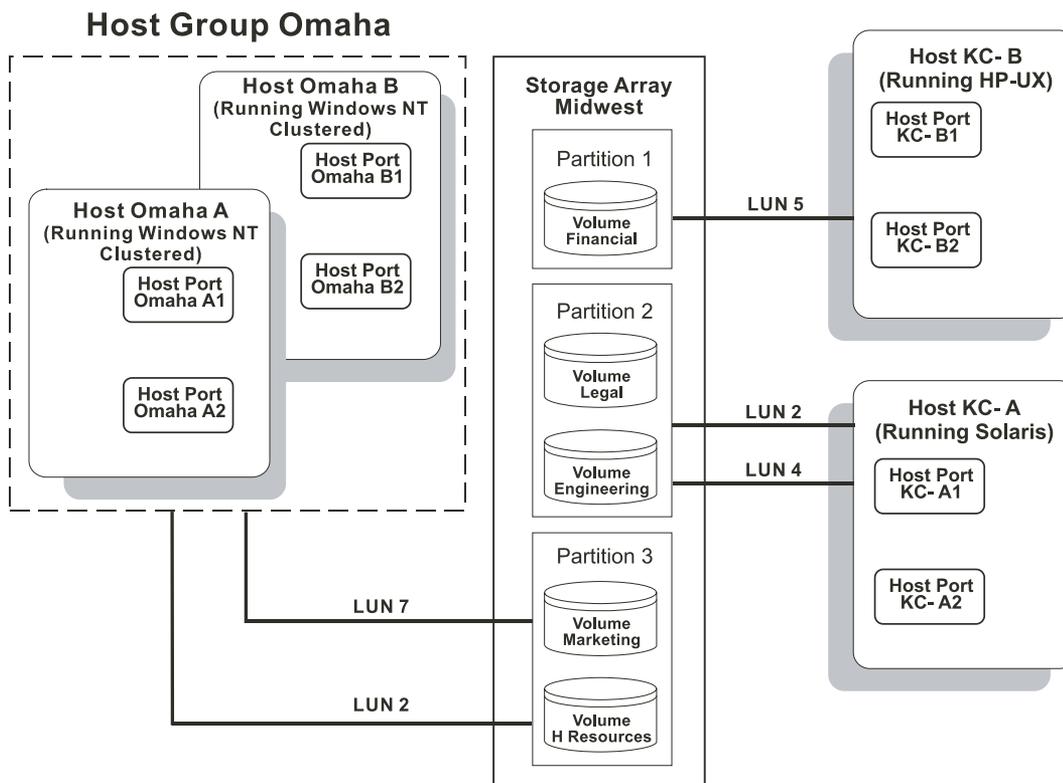


Figure 2-5 Heterogeneous Hosts Example

Maintaining Storage Arrays in the Management Domain

This section describes the maintenance of storage arrays in the management domain, which includes the following activities.

Use the **Enterprise Management** window to:

- Monitor the health of the storage arrays, as explained in the next section.
- Configure alert notifications using the event monitor (see “Event Monitor” on page 58 or the **Enterprise Management** window Help).

Note: To receive critical alerts, the **Enterprise Management** window must be running (it can be minimized).

Use the **Array Management** window to:

- Monitor the individual logical and physical components in an individual storage array (as explained in the **Array Management** window Help).
- Recover from storage array problems (as explained in “Recovering from Storage Array Problems” on page 65).
- Monitor storage array performance and tune performance, if necessary, (as explained in “Tuning Storage Arrays” on page 70).

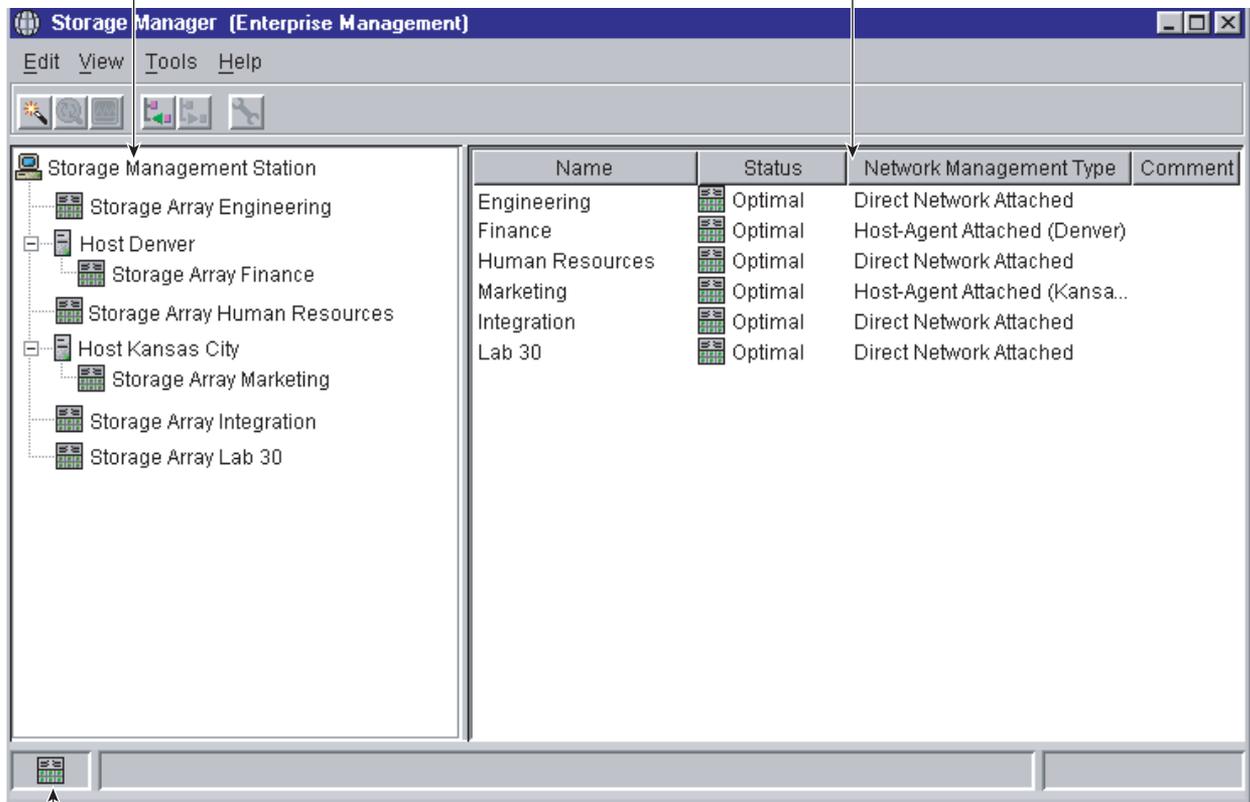
Monitoring the Health of Storage Arrays

Note: The **Enterprise Management** window or the event monitor must be running to receive notification of critical events for the storage arrays. In addition, alert notifications must be configured in the **Enterprise Management** window.

The **Enterprise Management** window provides a summary of the conditions of all known storage arrays in the management domain. Appropriate status indicators will be shown in the Device Tree, the Device Table, and in the health summary status area in the lower-left corner of the **Enterprise Management** window, as shown in Figure 2-6.

The Device Tree in this window provides a status for each storage array in the Management Domain. A status is shown for each network management connection.

The Device Table contains a status column that displays the status of each storage array



The Overall Health Status is a consolidated health status for all of the known storage arrays in the Management Domain

Figure 2-6 Monitoring the Health of Storage Arrays Using the Enterprise Management Window

Storage Array Status Icons

Table 2-6 provides information about the storage array status icons that display as follows:

- In the Device Tree, Device Table, and Overall Health Status areas of the **Enterprise Management** window.
- As the root node of the **Logical** view tree in the **Array Management** window.

Table 2-6 Storage Array Status Indications

Status	Icon	Indication
Optimal		Every component in the storage array is in the desired working condition.
Needs Attention		A problem with a storage array requires intervention to correct. To correct the problem, you should start the Array Management window for the particular storage array and then use Recovery Guru to pinpoint the cause of the problem and obtain appropriate instructions.
Fixing		A Needs Attention condition has been corrected and the storage array is currently transitioning to an Optimal state (for example, a reconstruction operation is in progress). A Fixing status requires no action unless you want to check on the progress of the operation in the Array Management window. Note: Some recovery actions cause the storage array state to change directly from Needs Attention to Optimal, without an interim state of Fixing. The Fixing status icon is not displayed in the Overall Health Status area (Optimal is displayed instead).

Table 2-6 Storage Array Status Indications (**continued**)

Status	Icon	Indication
Unresponsive		<p>The storage management station cannot communicate with the only controller or both controllers over its network management connection to the storage array.</p> <p>Note: The Unresponsive icon is not displayed in the Logical view of the Array Management window. If the Array Management window is open and the storage array becomes Unresponsive, the last known state (Optimal, Needs Attention, or Fixing) is shown.</p>
Contacting Device		<p>You have started the Enterprise Management window and the storage management software is establishing contact with the storage array.</p> <p>Note: This status is not displayed in the Logical view of the Array Management window.</p>

Event Monitor

The event monitor runs continuously in the background monitoring activity on a storage array and checking for critical problems (for example, impending drive failures or failed controllers). If the event monitor detects any critical problems, it can notify a remote system using e-mail and/or simple network management protocol (SNMP) trap messages whenever the **Enterprise Management** window is not running.

The event monitor is a separate program bundled with the client software and must be installed with the storage management software. The client/event monitor is installed on a storage management station or host connected to the storage arrays. For continuous monitoring, install the event monitor on a computer that runs 24 hours a day. Even if you choose not to install the event monitor, alert notifications must still be configured on the computer where the client software is installed, because alerts will be sent as long as the **Enterprise Management** window is running.

Figure 2-7 shows how the event monitor and the **Enterprise Management** window client software send alerts to a remote system. The storage management station contains a file with the name of the storage array being monitored and the address where alerts will be sent. The alerts and errors that occur on the storage array are continuously being monitored by the client software and the event monitor. The event monitor takes over for

the client after the client software package is shut down. When an event is detected, a notification is sent to the remote system.

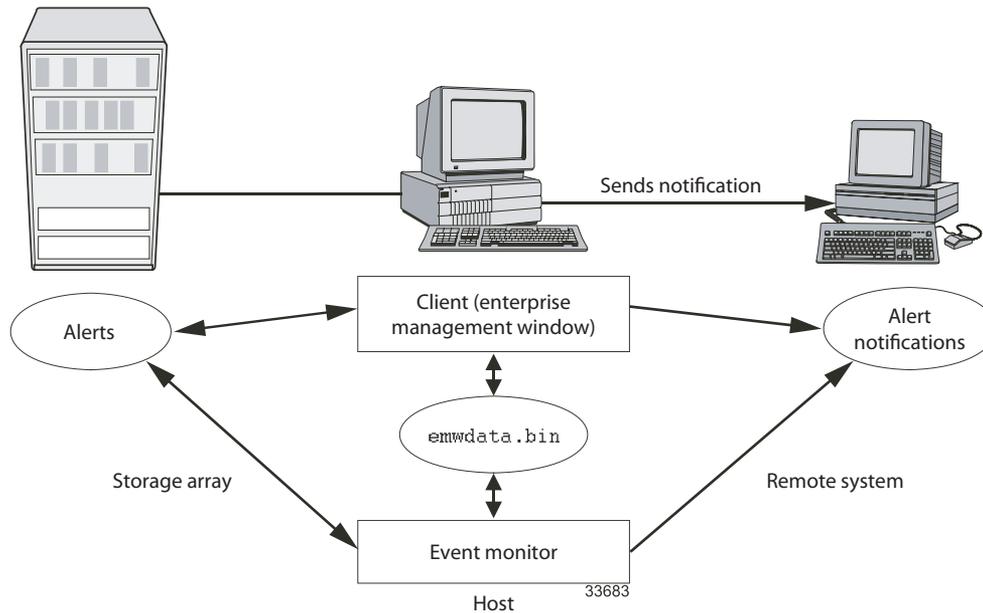


Figure 2-7 Event Monitor Configuration

Because the event monitor and the **Enterprise Management** window share the information to send alert messages, the **Enterprise Management** window contains some visual cues to assist in the event monitor installation and synchronization. The parts of the **Enterprise Management** window that are related to event monitoring are shown in Figure 2-7.

Using the event monitor involves three key steps:

1. Install the client software.

The event monitor is packaged with the client software and installs automatically with the client software. It is recommended that you run the event monitor on one machine that will run continuously. To prevent receipt of duplicate alert notifications of the same critical event on a storage array, disable the event monitor on all but one storage management station.

You must have administrative permissions to install software on the computer where the event monitor will reside. After the storage management software has been installed, the icon shown in Figure 2-8 will be present in the lower-left corner of the **Enterprise Management** window.

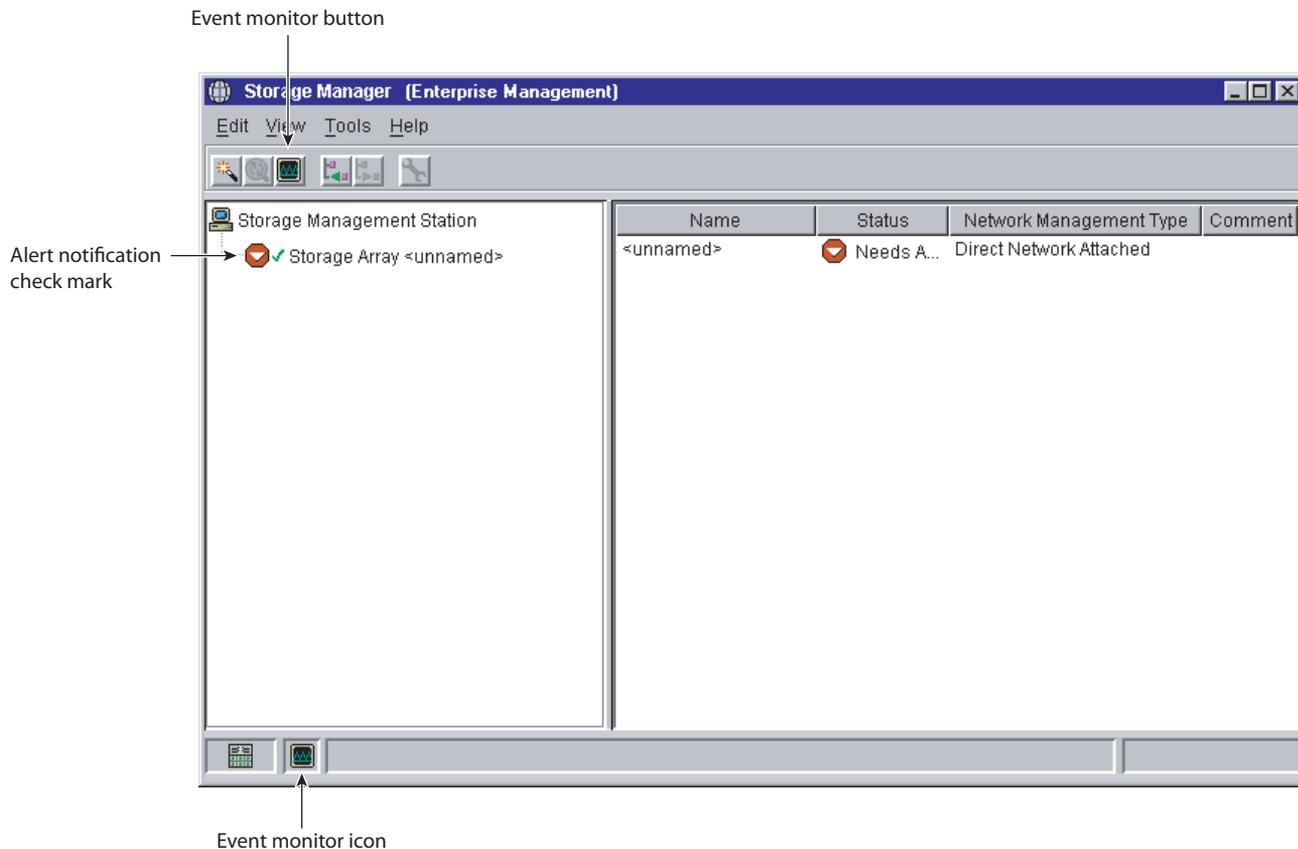


Figure 2-8 Event Monitor Example

2. Set up the alert destinations for the storage arrays you want to monitor from the **Enterprise Management** window.

A check mark indicates where the alert is set (storage management station, host, or storage array). When a critical problem occurs on the storage array, the event monitor will send a notification to the appropriate alert destinations that were specified.

3. Synchronize the **Enterprise Management** window and the event monitor. After it has been installed, the event monitor continues to monitor storage arrays and send alerts as long as it continues to run. If you make a configuration change in the **Enterprise Management** window, such as adding or removing a storage array or setting additional alert destinations, you should manually synchronize the **Enterprise Management** window and the event monitor.

Alert Notifications

Alert notification settings must be configured to receive e-mail or SNMP notifications if a critical event occurs on a storage array. The notification displays a summary of the critical event and details about the affected storage array, including:

- Name of the affected storage array
- Host IP address (only for a storage array managed through a host-agent)
- Host name/ID (shown as directly managed if the storage array is managed through each controller's Ethernet connection)
- Event error type related to an Event Log entry
- Date and time when the event occurred
- Brief description of the event

Note: To set up alert notifications using SNMP traps, you must copy and compile a management information base (MIB) file on the designated network management station.

Configuring Alert Notifications

1. Select a node in the **Enterprise Management** window that will display alert notifications for the storage arrays you want to monitor. You can set the alert notifications at any level:
 - Every storage array in the management domain
 - Every storage array attached and managed through a particular host
 - An individual storage array
2. Configure e-mail destinations, if desired. You must provide a mail server name and an e-mail sender address for the e-mail addresses to work.
3. Configure SNMP trap destinations, if desired. The SNMP trap destination is the IP address or the host name of a station running an SNMP service, such as a Network Management Station.

Customer Support Alert Notifications

Use the **Edit > Alert** option in the **Enterprise Management** window to configure the system to send e-mail notifications to a specified customer support group if a critical event occurs on one of the storage arrays. After it is properly set up, the e-mail alert notification includes a summary of the critical event, detailed information about the affected storage array, and custom contact information.

Note: If you do not configure this feature, the e-mail alert notification will contain only a summary of the critical event.

Configuring the Customer Support Alert Notification

1. Create a text file containing the contact information you want to send to the customer support group. For example, include the names and pager numbers of the network administrators.
2. Name the file `userdata.txt` and save it in the home directory (for example, `Winnt\profiles\ or /home/<name>/...`) on the client machine that is used to manage the storage array.

Note: This may be your host machine if you installed the client software on the host.

3. Configure the alert notifications by selecting the appropriate node in the Device Tree and then selecting **Edit > Alert > Destinations**. Specify the e-mail address of the customer support organization. For more information on selecting nodes for alert notifications, see “Configuring E-Mail and SNMP Alert Notification Settings” in the **Enterprise Management** window Help system.

Note: Remember to set up the mail server and sender address by selecting **Edit > Alert > Global Settings**.

Problem Notification

Note: The **Enterprise Management** window or the event monitor must be running (it can be minimized) for critical alert notifications to be sent. In addition, you must have configured the alert notifications in the **Enterprise Management** window.

Use the **Recovery Guru** to help troubleshoot storage array problems. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

When you monitor a storage array, there are several indicators of a storage array failure:

- A Needs Attention storage array icon is displayed in the following:
 - Overall Health Status area, Device Tree View, and Device Table of the **Enterprise Management** window.
 - **Logical** view of the **Array Management** window.
- The **Recovery Guru** toolbar button in the **Array Management** window changes from an Optimal to a Needs Attention status and flashes.
- Non-optimal component icons are displayed in the **Array Management** window **Logical/Physical View**.
- You receive critical SNMP or e-mail notifications.
- Hardware fault lights are displayed.

In Figure 2-9, the **Array Management** window for storage array Engineering shows indications of a components problem and a Needs Attention status in the **Logical/Physical View**.

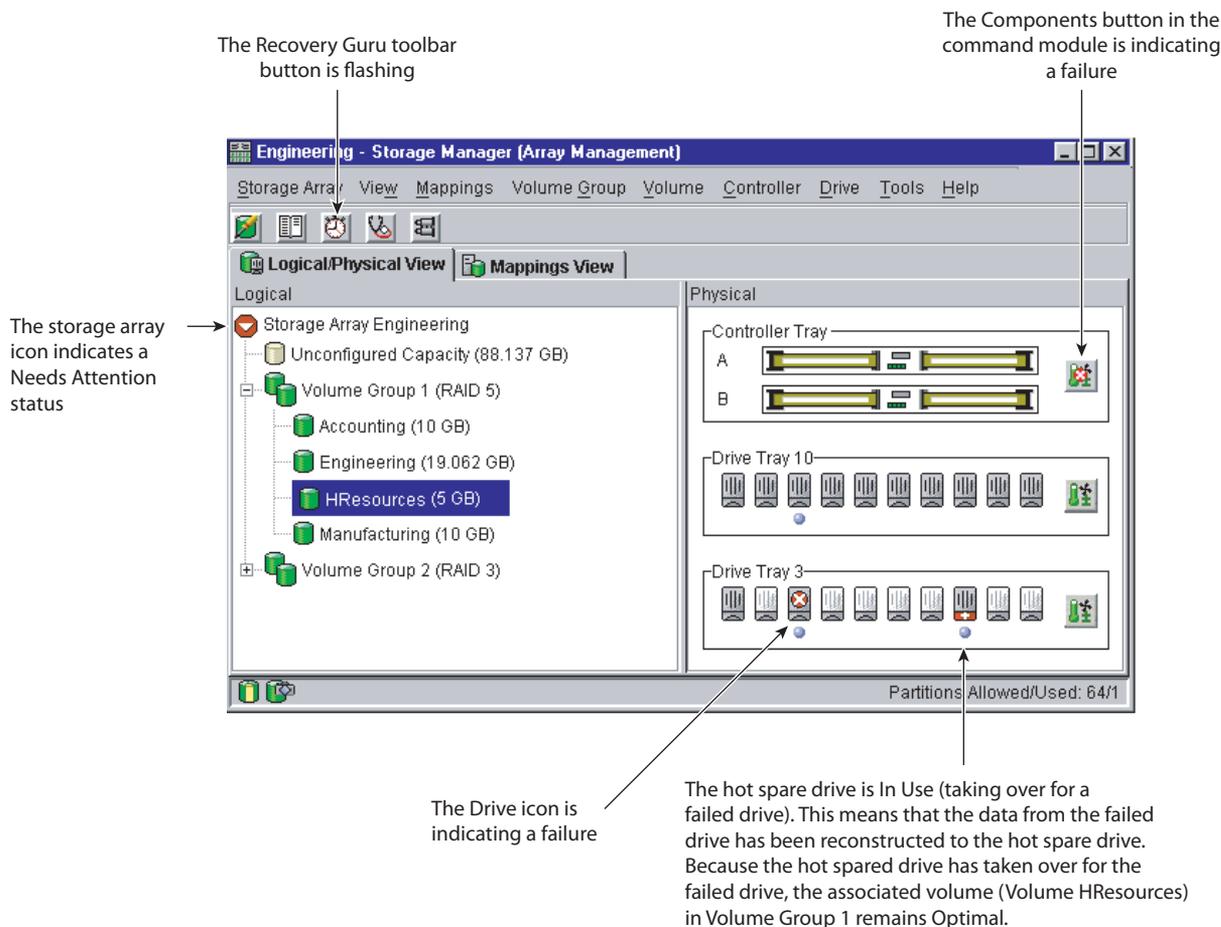


Figure 2-9 Problem Notification in the Array Management Window

You may also receive failure notifications about this storage array at a network storage management station or in e-mail. In addition, hardware fault lights display on the affected controller and drive enclosures.

Recovering from Storage Array Problems

The **Recovery Guru** is a component of the **Array Management** window in the ISSM-EE client package. The **Recovery Guru** diagnoses storage array problems and recommends recovery procedures to fix the problems. To display the **Recovery Guru**, select the **Recovery Guru** toolbar button in the **Array Management** window, called out in Figure 2-10, or select the **Storage Array > Recovery Guru** menu option.

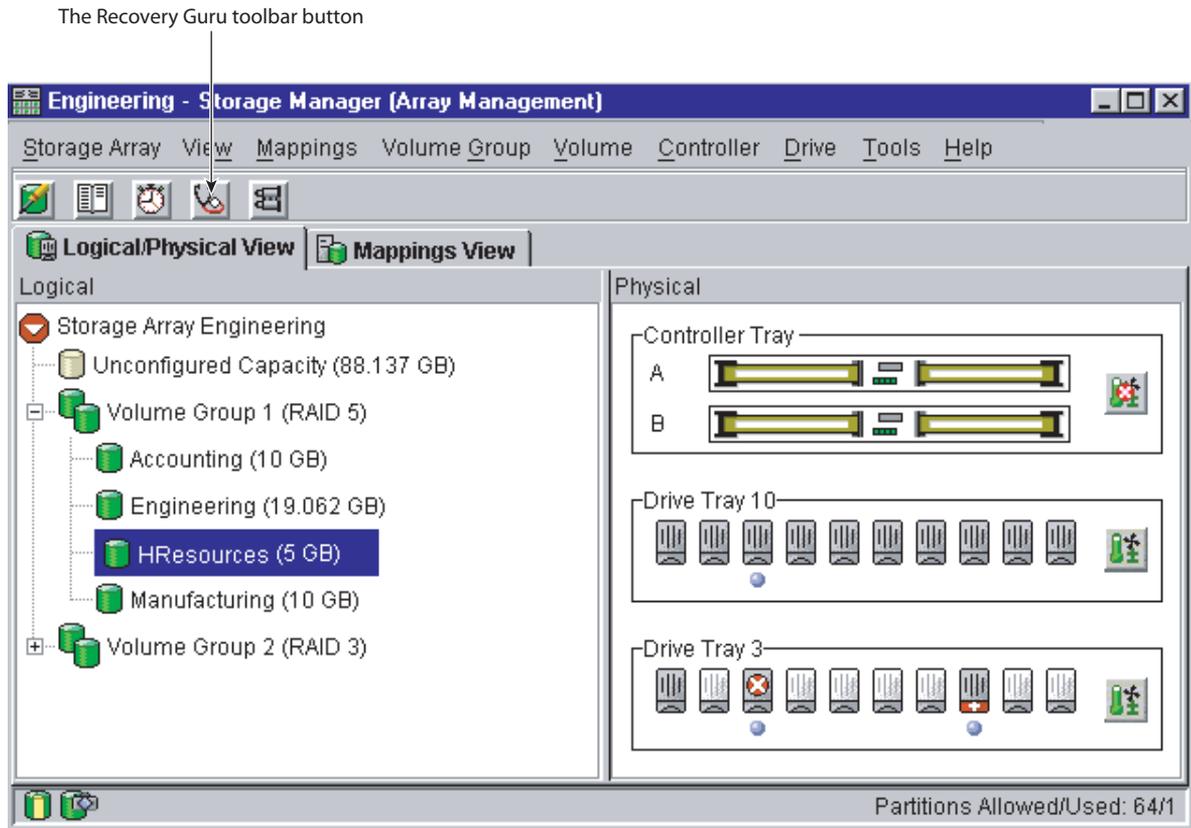


Figure 2-10 Recovery Guru Toolbar Button

Recovery Guru Example

The **Recovery Guru** window, as shown in Figure 2-11, is divided into three areas: **Summary**, **Details**, and **Recovery Procedure**. The **Summary** area lists storage array problems. The **Details** area displays information about the selected problem in the **Summary** area. The **Recovery Procedure** area lists the appropriate steps to follow for the selected problem in the **Summary** area.

For example, in Figure 2-11, the **Summary** area shows that there are two different failures in this storage array: a hot spare in use and a failed battery canister. The **Details** area shows that in Volume HResources, a hot spare drive in tray 10, slot 6 has replaced a failed drive in tray 3, slot 7. The **Recovery Procedure** area explains the cause of the problem selected in the **Summary** area, and describes the appropriate procedures needed to recover from this failure.

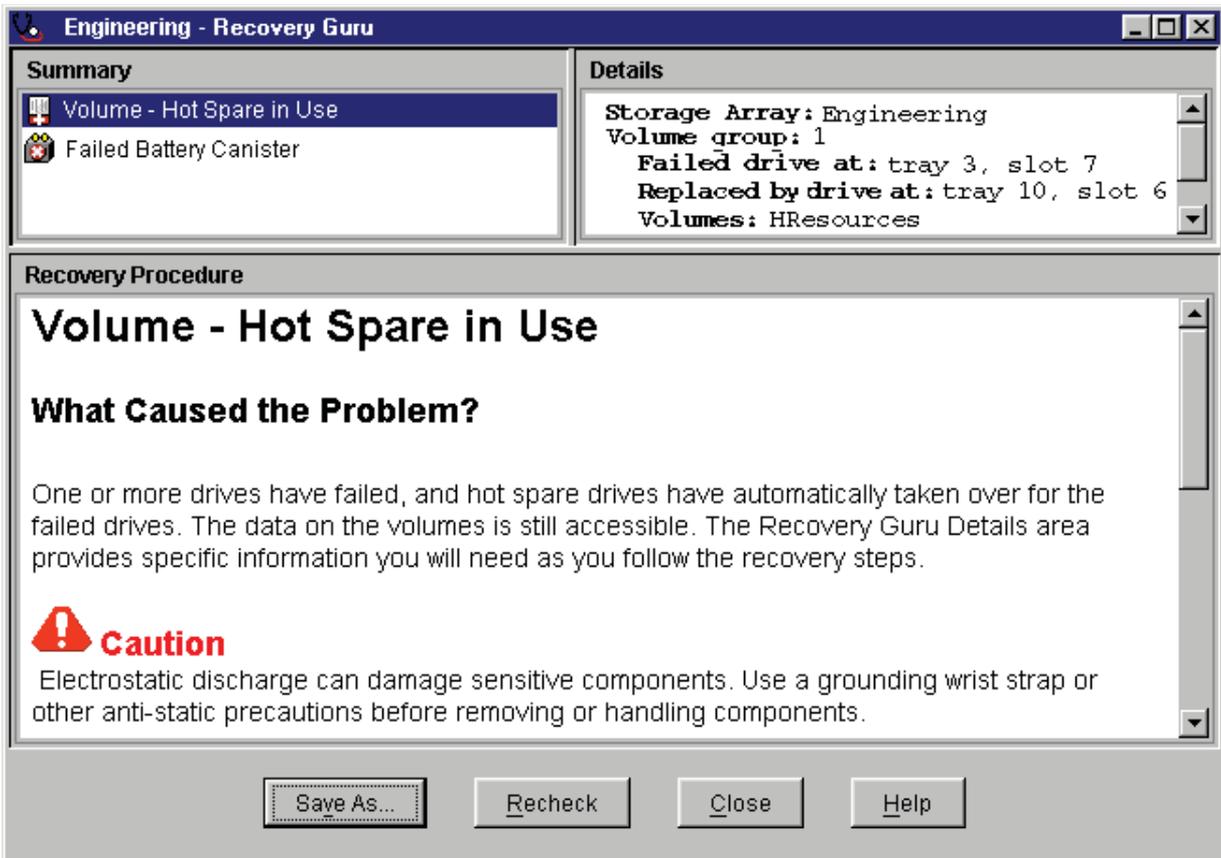


Figure 2-11 Recovery Guru Window

As you follow the recovery procedure to replace the failed drive, the storage array status changes to Fixing, the associated volume (HResources) status changes to Degraded-Modification in Progress, and the replaced drive status changes to Replaced. The data that was reconstructed to the hot spare drive is now being copied back to the replaced drive. These changes are shown in Figure 2-12.

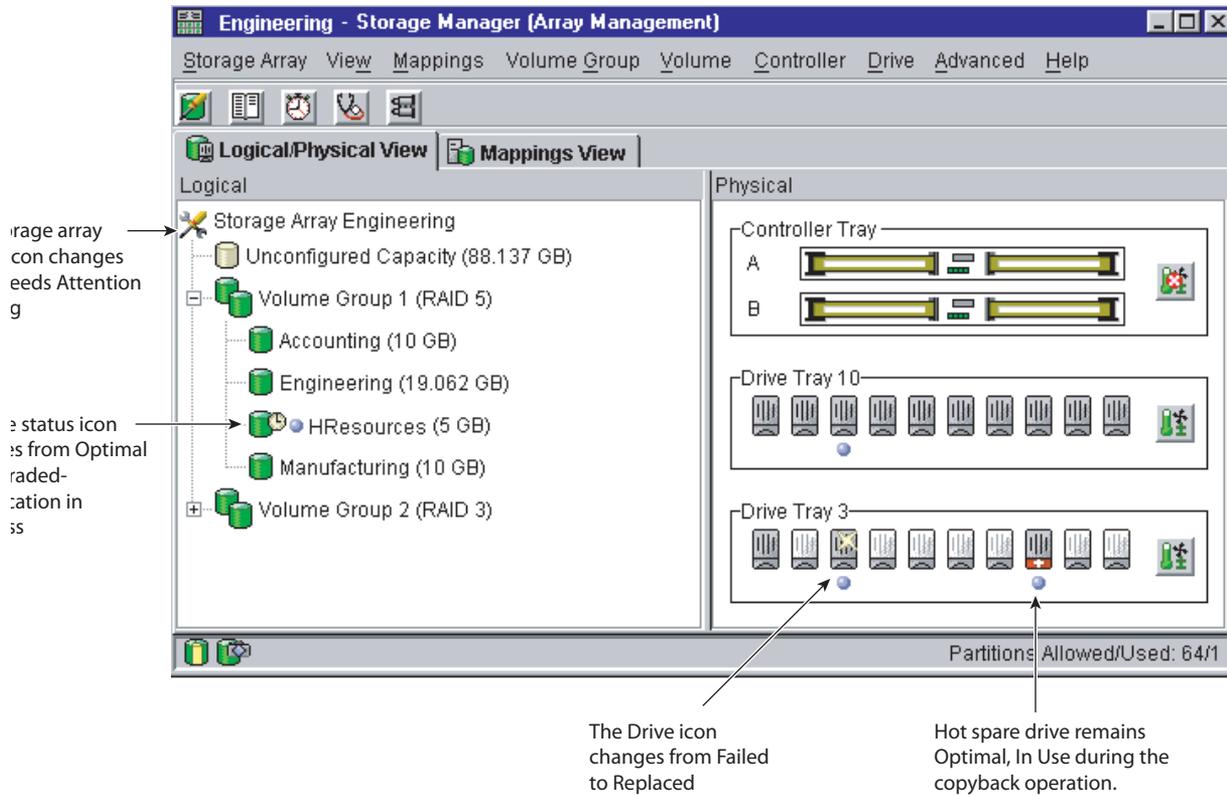


Figure 2-12 Status Changes During Example Recovery Operation

When the copyback operation is finished, the status icons change to reflect the optimal status of the components, as shown in Figure 2-13.

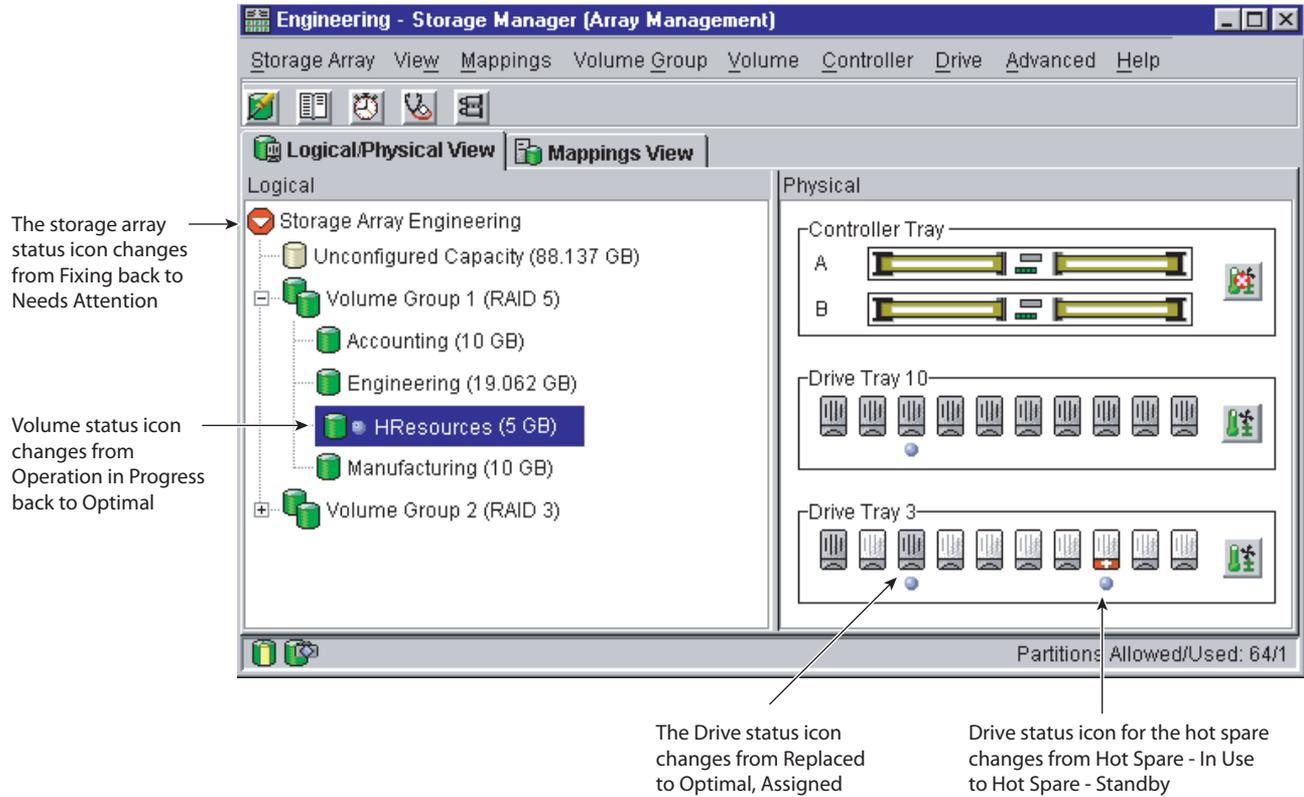


Figure 2-13 Status Changes When an Example Recovery Operation Completes

After you fix the failed drive in the drive enclosure, the following changes occur:

- The Storage Array icon in the **Logical** view returns to Optimal.
- The Storage Array icon in the **Enterprise Management** window changes from a Needs Attention status to an Optimal status.
- The **Recovery Guru** button stops blinking.

Note: For the **Recovery Guru** button to register Optimal status, the failed battery must also be replaced.

Tuning Storage Arrays

In general, when you are configuring a storage array, the default values should provide optimal performance. However, if you want to tune the storage array after initial configuration, several options are available in the **Array Management** window for optimizing storage array and application performance. These tuning options are described in Table 2-7.

Note: You can perform any of the tuning operations while the storage array is processing I/O.

Table 2-7 Tuning Options in Array Management Window

Option	Description
Storage Array > Monitor Performance	<p>Use this option to monitor storage array performance in real-time and save performance data to a file for later analysis. You can specify the volumes and/or controllers to monitor and the polling interval.</p> <p>Data provided from a session includes I/O load, cache hit percentage, how your I/O is allocated between reads and writes, transfer rates and I/O request rates for the volumes being monitored or for the entire storage array.</p> <ul style="list-style-type: none"> - The <i>cache hit percentage</i> is the percentage of reads that are serviced with data from the cache rather than requiring a read from disk. - A <i>transfer rate</i> is the amount of data in kilobytes per second that moves through the I/O data path (also called throughput). - An <i>I/O request rate</i> is the average number of I/O requests serviced per second. <p>For more information, see the Array Management window Help.</p>
Storage Array > Change Cache Settings	<p>Use this option to specify the start flush and stop flush parameters (expressed in percentage of cache memory that contains unwritten data) and the cache block size for the storage array.</p> <p>Low start and stop flush percentages increase data safety but decrease the cache hit percentage and therefore the I/O request rate.</p> <p>For more information, see “Protect Data in the Controller Cache Memory” on page 34 and the Array Management window Help.</p>

Table 2-7 Tuning Options in Array Management Window (**continued**)

Option	Description
Volume Group > Change RAID Level	<p>Use this option to change the RAID level of a volume group when either of the following occurs:</p> <ol style="list-style-type: none"> 1. The application read or write percentages on a volume indicate that its current RAID level may not be appropriate for desired performance. 2. You want to change the volume group's redundancy (such as changing from RAID 0 to RAID 5). <p>Applications with a high read percentage perform very well using RAID 5 volumes because of the outstanding read performance of the RAID 5 configuration.</p> <p>However, applications with a low read percentage (write-intensive) do not perform as well on RAID 5 volumes because of the way a controller writes data and redundancy data to the drives in a RAID 5 volume group. If there is a low percentage of read activity relative to write activity, you might consider changing the RAID level of a volume group from RAID 5 to RAID 1 for faster performance.</p>
Volume Group > Change Ownership	<p>If a particular controller has considerably more I/O activity than expected, consider moving a volume group to the other controller in the storage array by using this option.</p>
Volume > Change Segment Size	<p>A <i>segment</i> is the amount of data in kilobytes that the controller writes on a single drive in a volume before writing data on the next drive.</p> <p>When you create a volume, the default segment size is normally a good choice for the expected volume usage you specify. If volume usage is different than expected, use the Volume > Change Segment Size option to change the segment size of a volume.</p> <p>In a read-intensive application environment with random I/O access patterns, make sure the segment size is large enough to minimize the number of drives needed to satisfy an I/O request. This leaves other drives available to satisfy other requests.</p> <p>A <i>volume group data stripe</i> is the segment size multiplied by the number of drives in the volume group used for I/O. In a single-user environment with sequential I/O access patterns, performance is optimized when a single I/O request can be serviced with a single volume group data stripe. In this case, multiple disks are used for the same request, but each disk is only accessed once.</p>

Table 2-7 Tuning Options in Array Management Window (**continued**)

Option	Description
Volume > Properties	<p>Use this option to set following three tuning parameters:</p> <ul style="list-style-type: none"> - Write caching - Cache read-ahead - Volume Modification Priority <p><i>Write caching.</i> To maximize I/O request rates, enable write-caching on a volume. Write caching frees an application to continue after writing to the controller cache memory, a faster operation than a disk write. The cache memory is flushed to disk when it contains a percentage of unwritten data you specify using the Storage Array > Cache Settings option.</p> <p><i>Cache read-ahead.</i> If an individual volume is experiencing a low cache hit percentage, consider enabling cache read-ahead (or prefetch) for that volume. Cache read-ahead can increase the cache hit percentage for a sequential I/O workload. If cache read-ahead is enabled, the cache reads the data from the disk. But in addition to the requested data, the cache also fetches more data, usually from adjacent data blocks on the drive. This feature increases the chance that a future request for data could be fulfilled from the cache rather than requiring a disk access.</p> <p>The cache read-ahead multiplier values specify the multiplier to use for determining how many additional data blocks are read into cache. Choosing a higher cache read-ahead multiplier can increase the cache hit percentage for volumes with sequential I/O access patterns. (In the case of a random I/O pattern, the optimal multiplier is zero.)</p> <p><i>Volume modification priority.</i> This defines how much processing time is allocated for volume modification operations versus system performance. The higher the priority, the faster that volume modification operations are completed but the slower I/O is serviced. Volume modification operations include reconstruction, copyback, initialization, defragmentation, change of RAID level, and change of segment size.</p>

Table 2-7 Tuning Options in Array Management Window (**continued**)

Option	Description
Volume > Properties	You set the modification priority for each volume by using a slider bar on the Volume > Properties dialog. There are five relative settings on the reconstruction rate slider bar, ranging from Lowest to Highest . The actual speed of each setting is determined by the controller. Choose the Lowest setting to maximize the I/O request rate. If the controller is idle (not servicing any I/O), it will ignore the individual volume rate settings and process volume modification operations as quickly as possible.

Task Reference

There are two separate Help systems, one in the **Enterprise Management** window and one in the **Array Management** window. These are distinct Help systems that can only be accessed when you are in the appropriate window. Table 2-8 is a task-based index to the appropriate Help system.

Table 2-8 Task Reference: Populating a Management Domain

To perform this task...	Use this option	See the Help system in...
Adding a comment for a storage array	Edit > Comment	Enterprise Management window
Automatically discovering devices	Tools > Automatic Discovery	Enterprise Management window
Discovering newly attached host-agent managed storage arrays	Tools > Rescan	Enterprise Management window
Manually adding a device	Edit > Add Device	Enterprise Management window
Removing a device	Edit > Remove Device	Enterprise Management window

Table 2-9 describes the Storage Array Configuration Help system.

Table 2-9 Task Reference: Storage Array Configuration

To perform this task...	Use this option	See the Help system in...
Using the Script Editor	Tools > Execute Script	Enterprise Management window
Adding free capacity to a volume group	Volume Group > Add Free Capacity	Array Management window
Assigning/unassigning a hot spare drive	Drive > Assign Hot Spare and Drive > Deassign Hot Spare	Array Management window
Changing the enclosure order in the Physical view	Configure > enclosure Order	Array Management window
Consolidating free capacity on a volume group	Volume Group > Defragment	Array Management window
Creating a volume from free capacity	Configure > Create Volume	Array Management window
Creating a volume group and volumes from unconfigured capacity	Configure > Create Volume Group/Volume	Array Management window
Creating storage partitions	Configure > Topology/ Storage Partitioning	Array Management window
Downloading an NVSRAM configuration file	Storage Array > Download > NVSRAM	Array Management window
Downloading firmware to the controllers in a storage array	Storage Array > Download > Firmware	Array Management window
Increasing the unconfigured capacity of a storage array (deleting a volume group)	Volume Group > Delete	Array Management window
Increasing the free capacity of a storage array (deleting a volume)	Volume > Delete	Array Management window

Table 2-9 Task Reference: Storage Array Configuration (continued)

To perform this task...	Use this option	See the Help system in...
Performing an automatic configuration	Configure > Automatic Configuration	Array Management window
Resetting a configuration	Configure > Reset Configuration	Array Management window

Table 2-10 describes the Performance and Tuning Help system.

Table 2-10 Task Reference: Performance and Tuning

To perform this task...	Use this option	See the Help system in...
Changing cache settings on a storage array	Storage Array > Change Cache Settings	Array Management window
Changing cache settings on an individual volume	Volume > Properties	Array Management window
Changing media scan settings on a storage array	Storage Array > Change Media Scan Settings	Array Management window
Changing media scan settings on an individual volume	Volume > Properties	Array Management window
Changing the controller ownership of a volume group	Volume Group > Change Ownership	Array Management window
Changing the modification priority on an individual volume	Volume > Properties	Array Management window
Changing the RAID level of a volume group	Volume Group > Change RAID Level	Array Management window
Changing the segment size of a volume	Volume > Change Segment Size	Array Management window
Monitoring performance	Storage Array > Monitor Performance	Array Management window

Table 2-11 describes the Event Notification Help system.

Table 2-11 Task Reference: Event Notification

To perform this task...	Use this option	See the Help system in...
Interpreting an email or SNMP trap message	Not applicable	Enterprise Management window
Specifying and validating email and SNMP destinations for alert notifications	Edit > Alert > Destinations	Enterprise Management window
Specifying management domain global email alert settings	Edit > Alert > Global Settings	Enterprise Management window
Viewing and interpreting events stored in the Event Log	Storage Array > View Event Log	Array Management window

Table 2-12 describes the Recovering From Problems Help system.

Table 2-12 Task Reference: Recovering From Problems

To perform this task...	Use this option	See the Help system in...
Correcting a partially managed device	View > Partially Managed Devices	Enterprise Management window
Recovering a corrupted configuration file	Not applicable	Enterprise Management window
Recovering from an Unresponsive storage array	Not applicable	Enterprise Management window
Checking redundancy information on a volume group	Volume Group > Check Redundancy	Array Management window
Failing a drive	Drive > Fail	Array Management window
Identifying when to use Recovery Guru	Not applicable	Array Management window

Table 2-12 Task Reference: Recovering From Problems (continued)

To perform this task...	Use this option	See the Help system in...
Initializing drives, volumes, or volume groups	Drive > Initialize Volume > Initialize Volume Group > Initialize	Array Management window
Interpreting and using Recovery Guru information and procedures	Storage Array > Recovery Guru	Array Management window
Moving volume groups back to their preferred controller owners	Storage Array > Redistribute Volume Groups	Array Management window
Placing a controller in active or passive mode	Controller > Change Mode	Array Management window
Placing a controller online or offline	Controller > Place Online or Offline	Array Management window
Reconstructing a drive	Drive > Reconstruct	Array Management window
Reviving a drive	Drive > Revive	Array Management window

Table 2-13 describes the Miscellaneous System Administration Help system.

Table 2-13 Task Reference: Miscellaneous System Administration

To perform this task...	Use this option	See the Help system in...
Changing a password for a storage array	Storage Array > Change Password	Array Management window
Listing logical and physical components associated with a controller or drive	Controller > List Associated Components Drive > List Associated Components	Array Management window
Locating a storage array, volume group, or drive by flashing indicator lights	Storage Array > Locate Volume Group > Locate Drive > Locate	Array Management window
Renaming a storage array or volume	Storage Array > Rename Volume > Rename	Array Management window

Table 2-13 Task Reference: Miscellaneous System Administration (**continued**)

To perform this task...	Use this option	See the Help system in...
Synchronizing storage array controller clocks with the storage management station	Storage Array > Set Controller Clocks	Array Management window
Viewing a description of all components and properties of a storage array	Storage Array > Profile	Array Management window
Viewing volume, controller, or drive characteristics	Volume > Properties Controller > Properties Drive > Properties	Array Management window

Remote Volume Mirroring

This chapter includes the following sections:

- “Feature Concepts and Requirements” on page 79
- “Hardware and Software Installation” on page 108
- “Using the Remote Volume Mirroring Premium Feature” on page 134
- “Disaster Recovery” on page 158
- “Troubleshooting” on page 167

Feature Concepts and Requirements

This section introduces the following basic concepts and requirements of the Remote Volume Mirroring feature:

- “Introduction to Remote Volume Mirroring” on page 80
- “Data Replication” on page 83
- “Connectivity and Input/Output” on page 91
- “Using Other Premium Features and Remote Volume Mirroring” on page 94
- “Hardware and Software Requirements” on page 100
- “Recommended Configurations” on page 101
- “Hardware and Software Checklist” on page 106

Introduction to Remote Volume Mirroring

The Remote Volume Mirroring premium feature is used for online, real-time replication of data between storage arrays over a remote distance. In the event of a disaster or catastrophic failure at one storage array, you can promote a second storage array to take over responsibility for computing services. You must enable and activate the Remote Volume Mirroring premium feature on both the primary and the secondary storage arrays to create a mirror relationship.

This section introduces you to primary, secondary, and mirror repository volumes, and describes how they interact to replicate data between storage arrays using the Remote Volume Mirroring feature.

Primary and Secondary Volumes

When you create a remote volume mirror, a mirrored volume pair is created and consists of a primary volume at a primary storage array and a secondary volume at a secondary storage array. A standard volume may only be included in one mirrored volume pair.

The primary volume is the volume that accepts host I/O and stores application data. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the secondary volume. This process is known as a full synchronization and is directed by the controller owner of the primary volume. During a full synchronization, the primary volume remains fully accessible for all normal I/O operations.

The controller owner of the primary volume initiates remote writes to the secondary volume to keep the data on the two volumes synchronized. Whenever the data on the primary volume and the secondary volume becomes unsynchronized, the controller owner of the primary volume initiates a full synchronization.

The secondary volume maintains a mirror (or copy) of the data on its associated primary volume. The controller owner of the secondary volume receives remote writes from the primary volume controller owner, but will not accept host read or write requests. Hosts are able to read from the secondary volume, which appears as read-only.

The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure of the primary site, a role reversal can be performed to promote the secondary volume to a primary role. Hosts will then be able to access the newly promoted volume and business operations can continue.

Figure 3-1 shows the primary and secondary volumes displayed in the **Array Management** window for the primary storage array and the secondary storage array.

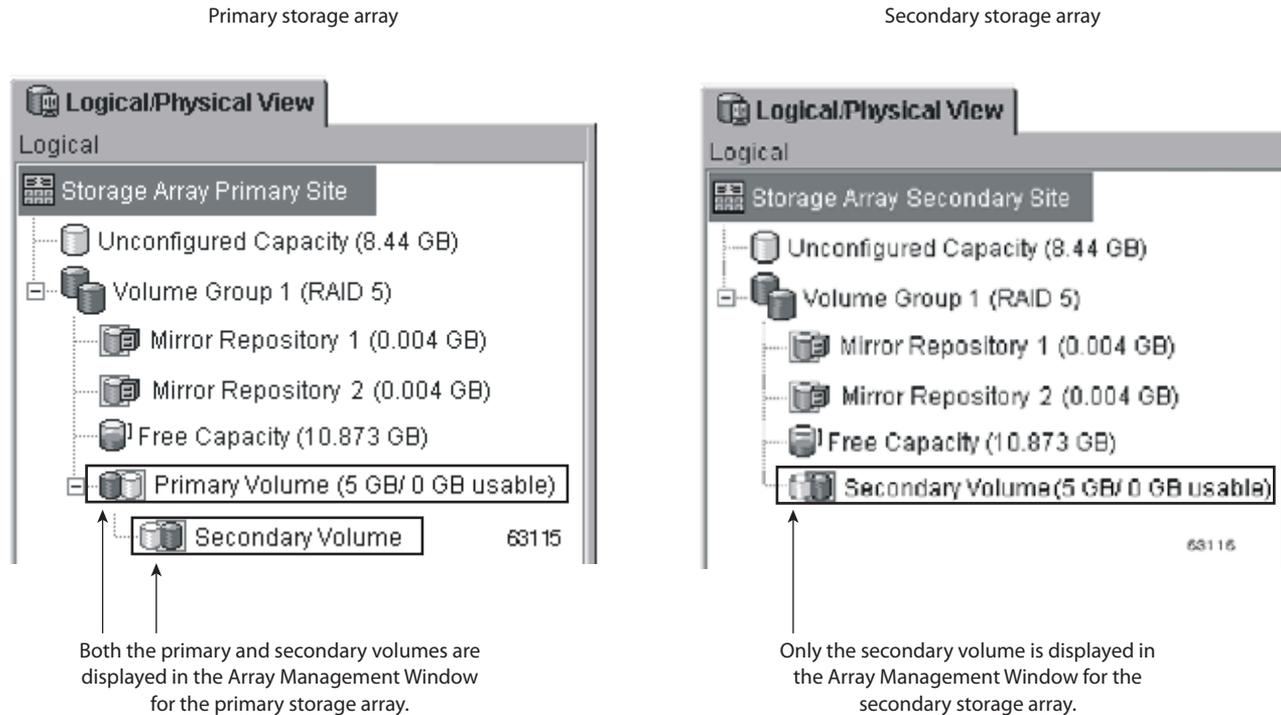


Figure 3-1 Primary and Secondary Volumes Displayed in the Array Management Window

Mirror Repository Volumes

A mirror repository volume is a special volume in the storage array created as a resource for the controller owner of the primary volume in a Remote Volume Mirror. The controller stores mirroring information on this volume, including information about remote writes that are not yet complete. The controller can use this information to recover from controller resets and accidental powering-down of storage arrays.

When you activate the Remote Volume Mirroring premium feature on the storage array, you create two mirror repository volumes, one for each controller in the storage array. An individual mirror repository volume is not needed for each Remote Volume Mirror.

When you create the mirror repository volumes, you specify the location of the volumes. You can either use existing free capacity or you can create a volume group for the volumes from unconfigured capacity and then specify the RAID level.

Figure 3-2 shows the mirror repository volumes displayed in the **Array Management** window for the primary storage array.

Because of the critical nature of the data being stored, the RAID level of mirror repository volumes must not be RAID 0 (for data striping). The required size of each volume is 128 MB, or 256 MB total for both mirror repository volumes of a dual-controller storage array. In the previous versions of the Remote Volume Mirroring feature, the mirror-repository volumes required less disk drive space, but now you must upgrade them to use the maximum number of mirror relationships.

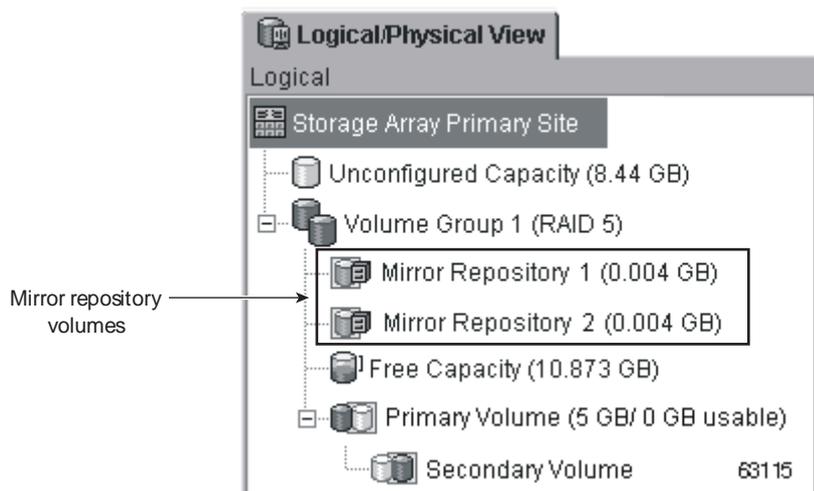


Figure 3-2 Mirror Repository Volumes Displayed in the Array Management Window

Mirror Relationships

Prior to creating a mirror relationship, the Remote Volume Mirroring premium feature must be enabled on both the primary and secondary storage arrays. A secondary volume candidate must be created on the secondary site if one does not already exist and must be a standard volume of equal or greater capacity than the associated primary volume.

When secondary volume candidates are available, a mirror relationship can be established in the storage management software by identifying the storage array containing the primary volume and the storage array containing the secondary volume.

When the mirror relationship is first created, a full synchronization automatically occurs, with data from the primary volume copied in its entirety to the secondary volume. For more information on creating mirror relationships, see “Creating Mirror Relationships” on page 142.

Data Replication

Data replication between the primary volume and the secondary volume is managed by the controllers and is transparent to host machines and applications. This section describes how data is replicated between storage arrays participating in Remote Volume Mirroring and the actions taken by the controller owner of the primary volume if a link interruption occurs between storage arrays.

Write Modes

When the controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to a mirror repository volume. Then the controller writes the data to the primary volume. The controller then starts a remote write operation to copy the affected data blocks to the secondary volume on the secondary storage array.

Three write mode options affect when the I/O completion indication is sent back to the host: Synchronous, Write Consistency, and Asynchronous.

Synchronous Mode

In the event of a communication failure, the Synchronous write mode offers the best chance of full data recovery from the secondary storage array. However, this action is often accomplished at the expense of host I/O performance. When you select this write mode, any host write requests are written to the primary volume and then are copied to the secondary volume.

After the host write request has been written to the primary volume and the data has been successfully copied to the secondary volume, the controller removes the log record on the mirror repository volume and sends an I/O completion indication back to the host system. The Synchronous write mode is selected by default and is the recommended write mode.

Because the controller does not send the I/O completion to the host until the data has been copied to both the primary and secondary volumes, this mirroring operation is called synchronous. When a read request is received from a host system, the controller owner of the primary volume handles the request normally and no communication takes place between the primary and secondary storage arrays. See Figure 3-3 for a graphical representation of the data replication process.

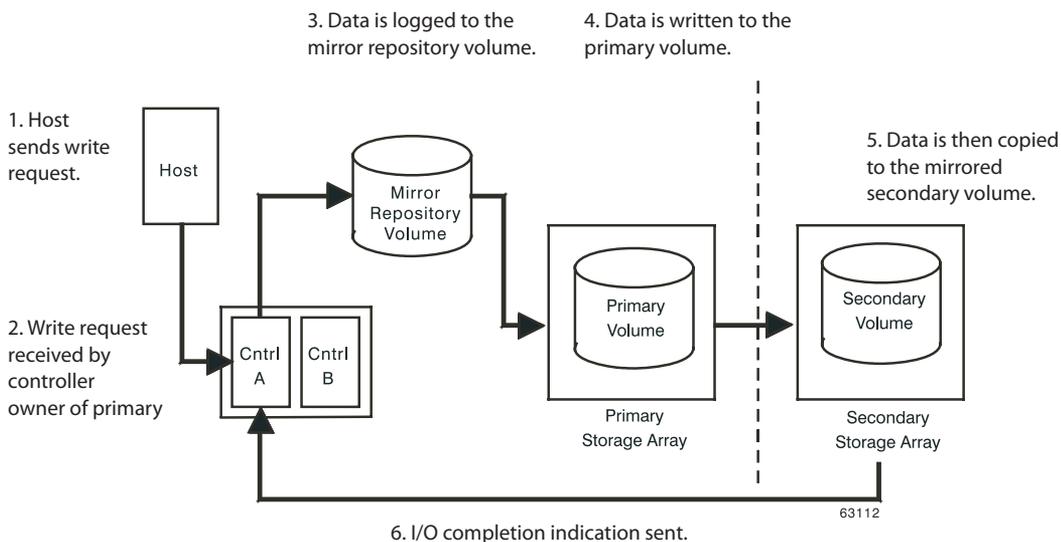


Figure 3-3 Data Replication between Storage Arrays

Write Consistency Mode

Write Consistency write mode preserves the write order when mirroring data from the primary volume to the secondary volume. A write consistency group is an interdependent group that exists when multiple mirror relationships exist on a single storage array. In addition, a write consistency group is created to preserve write consistency. The write consistency group preserves the write order sent by the application to multiple volumes. All mirror relationships in the write consistency group maintain the same order when sending writes from the primary volumes to their corresponding secondary volumes.

When you select Write Consistency write mode for multiple mirror relationships on the same storage array, the order that the data is synchronized is preserved. For example, the write from Mirror 1 is always performed before the write from Mirror 2. Selecting Write Consistency write mode for a single mirror relationship does not change the process in which data is replicated. More than one mirror relationship must reside on the primary storage array for the replication process to change.

Note: Only one write consistency group is supported. All mirrors created in Write Consistency write mode are in the same group by default.

The data on the secondary, remote storage array cannot be considered fully synchronized until all Remote Volume Mirrors in the write consistency group are synchronized. If one mirror relationship in a write consistency group becomes unsynchronized, all of the mirror relationships in the write consistency group will become unsynchronized. Any write activity to the remote, secondary storage arrays will be prevented to protect the consistency of the remote data set.

For example, Remove Volume Mirroring is configured between two storage arrays using the campus configuration. At the primary site, the primary storage array has three defined mirror relationships (RVM-A, RVM-B, and RVM-C). Each mirror relationship is configured to copy data to the secondary storage array. All three mirror relationships are configured to use Write Consistency to preserve write order. If the mirrored pair RVM-A becomes unsynchronized due to a link interruption, the controller automatically transitions RVM-B and RVM-C into an Unsynchronized state. This Unsynchronized state remains until communication can be resumed.

Note: When you issue a suspend command or resume command on a write consistency group, you suspend or resume the entire group, not just a single mirror. Wait until the entire group has finished the synchronization before you let applications send new I/O. This action makes sure that the write order of the new I/O is correct.

Asynchronous Write Mode

The Asynchronous write mode offers faster host I/O performance, but it does not guarantee that the copy has been successfully completed before you process the next write request. When you select this write mode, host write requests are written to the primary volume. Also, the controller sends an I/O completion indication back to the host system without consideration of when the data was successfully copied to the secondary storage array.

Resynchronization Methods

The controllers manage data replication between the primary volume and the secondary volume in a mirror relationship. Data replication is transparent to host machines and applications. When the controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to a mirror repository volume. The controller then writes the data to the primary volume. The controller then starts a write operation to copy the affected data to the secondary volume on the secondary storage system.

If a link interruption or a volume error prevents communication with the secondary storage system, the controller owner of the primary volume transitions the mirrored pair into an Unsynchronized status. The controller owner also sends an I/O completion to the host sending the write request. The host can continue to issue write requests to the primary volume, but remote writes to the secondary volume do not take place. Refer to Figure 3-4.

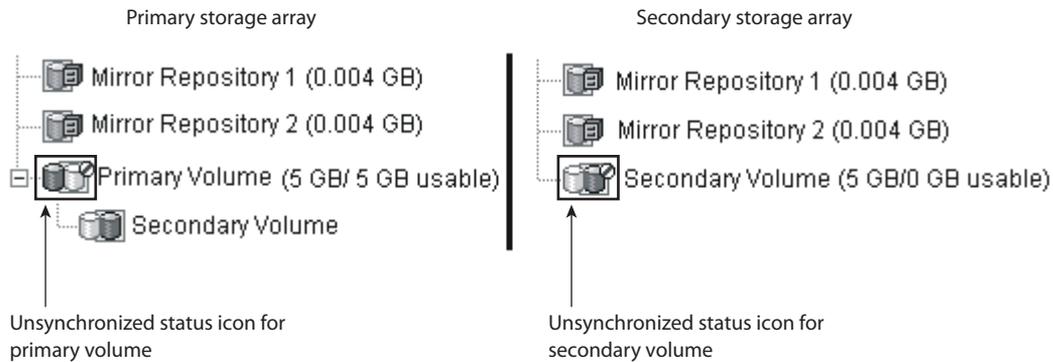


Figure 3-4 Mirrored Volume Pair in Unsynchronized Status

When connectivity is restored between the controller owner of the primary volume and the controller owner of the secondary volume, a full synchronization will automatically take place. The mirrored pair will then transition from an Unsynchronized status to a Synchronization-in-Progress status, as shown in Figure 3-5.

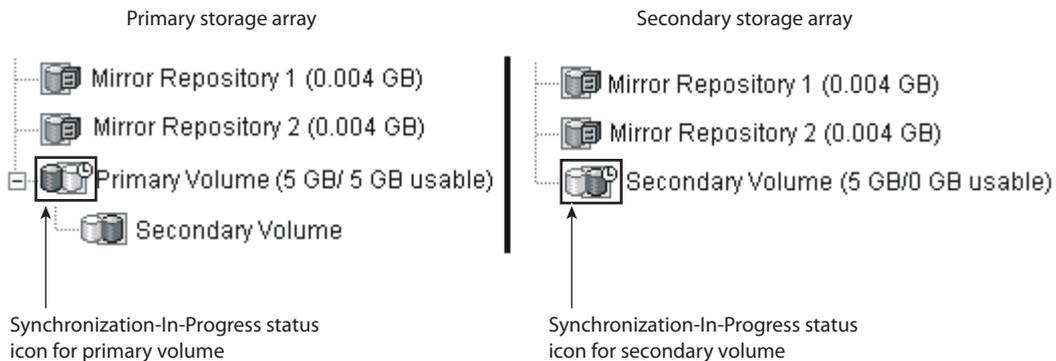


Figure 3-5 Mirrored Volume Pair in Synchronization-in-Progress Status

Important: Possible loss of data access – Any communication disruptions between the primary storage system and the secondary storage system while resynchronization is underway could result in a mix of new data and old data on the secondary volume. This action would render the data unusable in a disaster-recovery situation.

The primary controller will also mark the mirrored pair as Unsynchronized when a volume error on the secondary side prevents the remote write from completing. For example, an offline or a failed secondary volume can cause the Remote Volume Mirror to become Unsynchronized. When the volume error is corrected (the secondary volume is placed online or recovered to an Optimal status), then a full synchronization automatically begins, and the mirrored pair transitions to a Synchronization-in-Progress status. For more information on Remote Volume Mirroring statuses, see “Using the Remote Volume Mirroring Premium Feature” on page 134.

Two resynchronization methods are available in the current release of storage management software: Manual Resynchronization, which is the recommended method, and Automatic Resynchronization.

Manual Resynchronization

When enabled, you can manually start resynchronization of the data on the primary volume and the secondary volume after communication has been restored to the unsynchronized mirrored pair.

Important: Only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume.

If you select the Manual Resynchronization option, you can manage the resynchronization process in a way that provides the best opportunity for data recovery. When you select this option and a communication failure occurs between the primary volumes and the secondary volumes, the mirror relationship transitions to an Unsynchronized state. Any write requests to the primary volume are logged. A Needs Attention status appears for the storage system.

After the controller owner of the primary volume detects that communication has been restored, the mirror relationship remains in an Unsynchronized state until you select the **Volume >> Remote Volume Mirroring >> Resume** option in the storage management software.

Automatic Resynchronization

The controller owner of the primary volume resynchronizes the volumes when the Automatic Resynchronization option is selected. Resynchronization also happens when communication failure occurs between the primary storage system and the secondary storage system. The primary volume and the secondary volume are resynchronized immediately after detecting that communication has been restored.

Important: Only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume.

General Performance Considerations

The following general performance considerations should be noted when creating Remote Volume Mirrors.

- The controller owner of a primary volume performs a full synchronization in the background while processing local I/O writes to the primary volume and associated remote writes to the secondary volume. Because the full synchronization diverts controller processing resources from I/O activity, it can have a performance impact to the host application.

To reduce the performance impact, you can set the Synchronization Priority Level to determine how the controller owner will prioritize the full synchronization relative to other I/O activity. The following are some basic guidelines to setting the Synchronization Priority Level.

- A full synchronization at the Lowest Synchronization Priority Level will take approximately eight times as long as a full synchronization at the Highest Synchronization Priority Level.
- A full synchronization at the Low Synchronization Priority Level will take approximately six times as long as a full synchronization at the Highest Synchronization Priority Level.
- A full synchronization at the Medium Synchronization Priority Level will take approximately three and a half times as long as a full synchronization at the Highest Synchronization Priority Level.
- A full synchronization at the High Synchronization Priority Level will take approximately twice as long as a full synchronization at the Highest Synchronization Priority Level.

For more information on setting the Synchronization Priority Level, see “Using the Remote Volume Mirroring Premium Feature” on page 134.

- When the mirrored volume pair is in a Synchronization-in-Progress status, all host write data is copied to the remote system in processing the I/O. Both controller I/O bandwidth and I/O latency can affect host write performance. Host read performance is not affected by the mirroring relationship.
- The time that it takes for data to be copied from the primary volume to the secondary volume may impact overall performance and is primarily because of the delay and overhead of copying data to the Remote Volume Mirror. Some delay may also occur because of the limit to the number of simultaneous writes.

Link Interruptions or Secondary Volume Errors

When processing write requests, the primary controller might be able to write to the primary volume, but a link interruption prevents communication with the remote, secondary controller.

In this case, the remote write cannot complete to the secondary volume, and the primary volume and the secondary volume are no longer mirrored correctly. The primary controller transitions the mirrored volume pair into an Unsynchronized status and sends an I/O completion to the primary host. The primary host can continue to write to the primary volume, but remote writes cannot take place.

When connectivity is restored between the controller owner of the primary volume and the controller owner of the secondary volume, a resynchronization takes place. This resynchronization can take place either automatically or manually. The type of resynchronization depends on which method you chose when setting up the mirror relationship. During the resynchronization, only the blocks of data that have changed on the primary volume during the link interruption are copied to the secondary volume. After the resynchronization begins, the mirrored volume pair transitions from an Unsynchronized status to a Synchronization-in-Progress status.

The primary controller also marks the mirrored volume pair as unsynchronized when a volume error on the secondary side prevents the remote write from completing. For example, an offline or a failed secondary volume can cause the remote mirror operation to become unsynchronized. When the volume error is corrected (the secondary volume is placed online or recovered to an Optimal status), a full synchronization (automatic or manual) is required. The mirrored pair then transitions to a Synchronization-in-Progress status.

Connectivity and Input/Output

The Remote Volume Mirroring feature requires a dedicated host port for mirroring data between storage arrays. This section describes the typical configurations used to connect storage arrays for Remote Volume Mirroring, controller ownership of volumes within storage arrays, the maximum storage array distance permitted, and some general performance considerations.

Once the Remote Volume Mirroring feature has been activated, one Fibre Channel host-side I/O port on each controller is solely dedicated to mirroring operations. For example, in the primary storage array, controller ports A2 and B2 are dedicated to mirroring operations. In the secondary storage array, controller ports A2 and B2 are also dedicated to mirroring operations.

In IS4500 controller enclosures, the controller ports that are dedicated to mirroring operations in both the primary storage system and the secondary storage system are A4 and B4, respectively. Any host-initiated I/O operations are not accepted by the dedicated port. Any requests received on this dedicated port are accepted only from another controller that is participating in the mirror relationship.

Any host-initiated I/O operations will not be accepted by the dedicated port, and any requests received on this dedicated port will only be accepted from another controller participating in the mirror relationship.

Controller ports dedicated to Remote Volume Mirroring must be attached to a Fibre Channel fabric environment with support for the Directory Service and Name Service interfaces.

Fibre Channel Fabric Configurations

Three Fibre Channel fabric configurations are supported in this version of the storage management software. The level of redundancy available will depend on the type of configuration selected.

- **Highest Availability Campus Configuration** – A fully redundant configuration, which consists of a primary storage array and a secondary storage array. For more information, see “Highest Availability Campus Configuration” on page 101.
- **Campus Configuration** – A lower-cost configuration, which consists of a primary storage array and a secondary storage array. For more information, see “Campus Configuration” on page 104.

- Intra-Site Configuration – Lowest-cost configuration, where the primary storage array and secondary storage array are located within the same site. For more information, see “Intra-Site Configuration” on page 105.

For instructions on how to configure storage arrays, see “Hardware and Software Installation” on page 108.

Volume Ownership

The controller owner of the primary volume will only attempt to communicate with its mirrored controller in the secondary storage array. As shown in Figure 3-6, Controller A in the primary storage array only attempts communication with Controller A in the secondary storage array.

The controller (A or B) that owns the primary volume determines the controller owner of the secondary volume. If the primary volume is owned by Controller A on the primary storage array, the secondary volume is therefore owned by Controller A on the secondary storage array. If primary Controller A cannot communicate with secondary Controller A, no controller ownership changes take place.

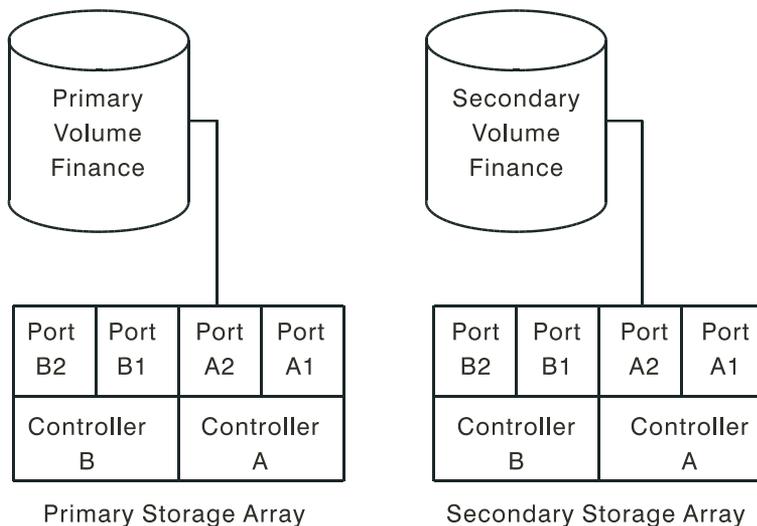


Figure 3-6 Controller A as Controller Owner

When an I/O path error causes a volume ownership change on the primary storage array, or if the storage administrator changes the controller owner of the primary volume, the next remote write processed will automatically trigger an ownership change on the secondary storage array.

As shown in Figure 3-7, if a primary volume is owned by Controller A and the controller owner is changed to Controller B, the next remote write will change the controller owner of the secondary volume from Controller A to Controller B.

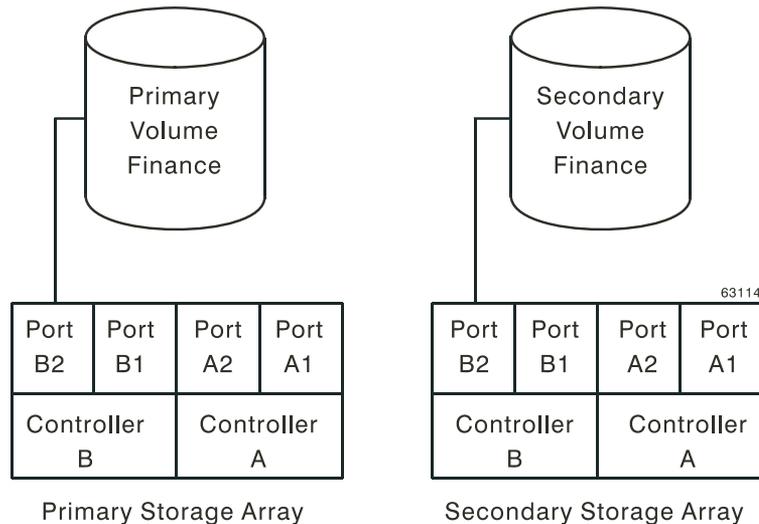


Figure 3-7 Controller B as Controller Owner

Because controller ownership changes on the secondary storage array are controlled by the primary controller, they do not require any special intervention and cannot be manually changed by the storage administrator.

Distance Between Primary and Secondary Storage Array

The maximum distance permitted between storage arrays participating in a mirror relationship is governed by the distance limits of the Fibre Channel links. Using standard Single-Mode Fiber technology, a maximum link distance of 10 km (6.25 miles) between storage arrays can be achieved.

Note: Specialized third-party products, such as Fibre Channel over IP (FCIP) switches, support distances of up to approximately 4000 km (2500 miles) between storage systems in a Fibre Channel environment.

Using Other Premium Features and Remote Volume Mirroring

This section describes how the Remote Volume Mirroring premium feature can be used in conjunction with other premium features available in this version of the storage management software.

Storage Partitioning

Storage Partitioning is a premium feature that enables hosts to share access to volumes in a storage array. A storage partition is created when you define a collection of hosts (a host group) or a single host and then define a volume-to-logical unit number (LUN) mapping. This mapping allows you to define what host group or host will have access to a particular volume in the storage array.

The storage partition definitions for the primary and secondary storage arrays are independent of each other. If these definitions are put in place while the volume is in a secondary role, it will reduce the administrative effort associated with site recovery if it becomes necessary to promote the volume to a primary role. For more information on defining storage partitions, see the **Array Management** window online help.

Figure 3-8 shows the **Mappings View** of the **Array Management** window, where Primary Volume Accounting has been included in a partition accessible by Host Group Campus East, and Secondary Volume Accounting 2 has been included in a partition accessible by Host Group Campus West.

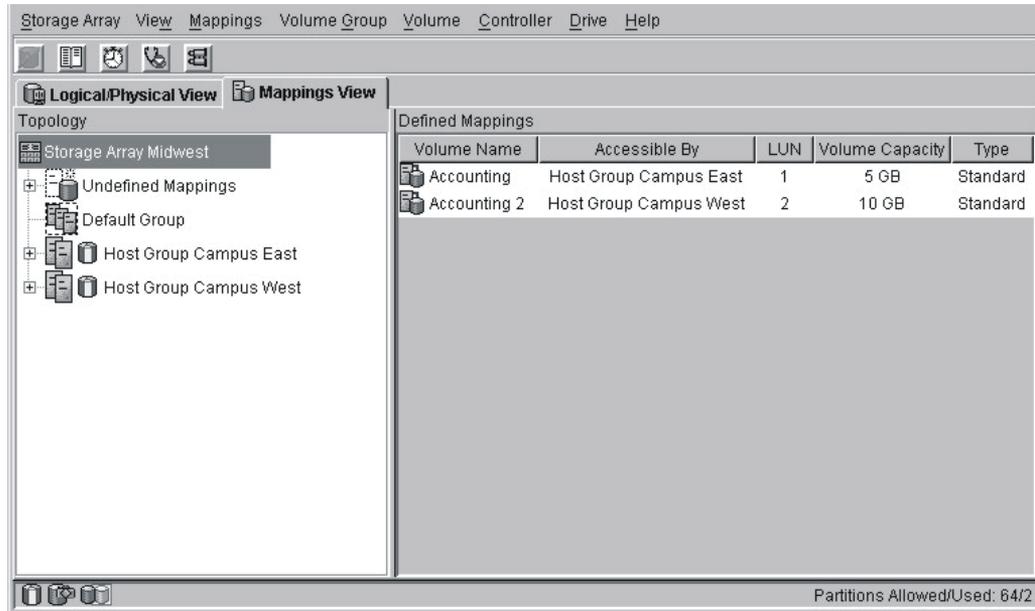


Figure 3-8 Storage Partitioning Example

Snapshot Volume

A snapshot volume is a point-in-time image of a volume and is typically created so that an application, such as a backup application, can access the snapshot volume and read the data while the base volume remains online and accessible to hosts.

The volume for which the point-in-time image is created is known as the base volume and must be a standard volume in the storage array. For the Remote Volume Mirroring feature, the base volume of a snapshot volume is permitted to be a candidate for the primary volume in a mirror relationship. The snapshot repository volume stores information about all data altered since the snapshot was created. Figure 3-9 shows the primary volume (also the base volume for the snapshot volume), the secondary volume, and the snapshot volume in the **Array Management** window for the primary storage array.

Note: Do not mount the snapshot volume on the same server as the primary volume. Also, you cannot gain “read-only” access to a remote volume if you are using a journaling file system. A journaling file system does not let you mount the remote volume in Windows (NTFS); however, you can mount the snapshot of the remote volume.

Secondary volumes are restricted from being the base volume for a snapshot volume because any full synchronization process would fill the snapshot repository volume, resulting in a failed snapshot volume.

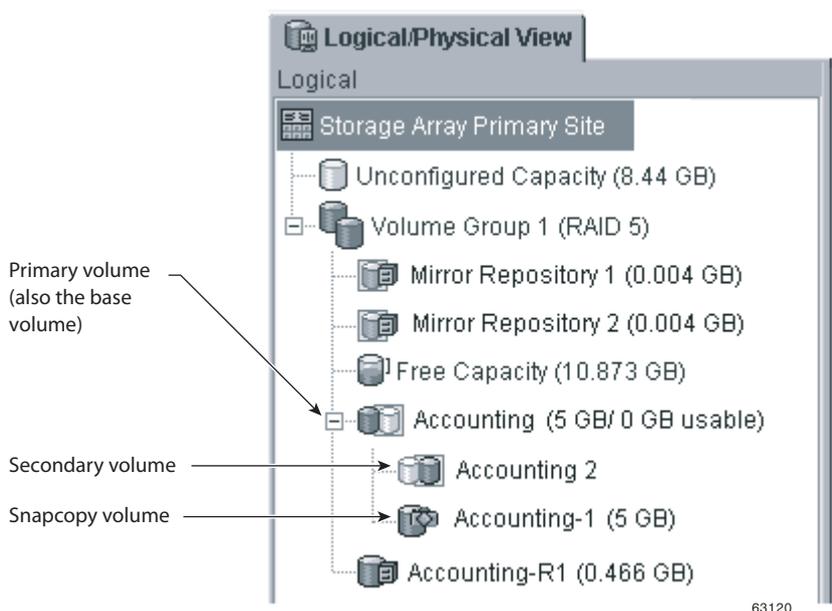


Figure 3-9 Snapshot Volume Included in a Mirror Relationship

Snapshot Volume Failure

In this version of the storage management software, snapshot volumes based on the primary volume of a mirrored volume pair are permitted.

If a role reversal is performed and the primary volume is demoted to the secondary volume role, all snapshot volumes associated with the demoted primary volume will be failed.

Note: A secondary volume **cannot** be selected to be the base volume of a snapshot volume.

As shown in Figure 3-10, Primary Volume Accounting and Secondary Volume Accounting 2 are displayed in the **Logical** view of the **Array Management** window for the primary storage array. A snapshot volume of Primary Volume Accounting has been created.

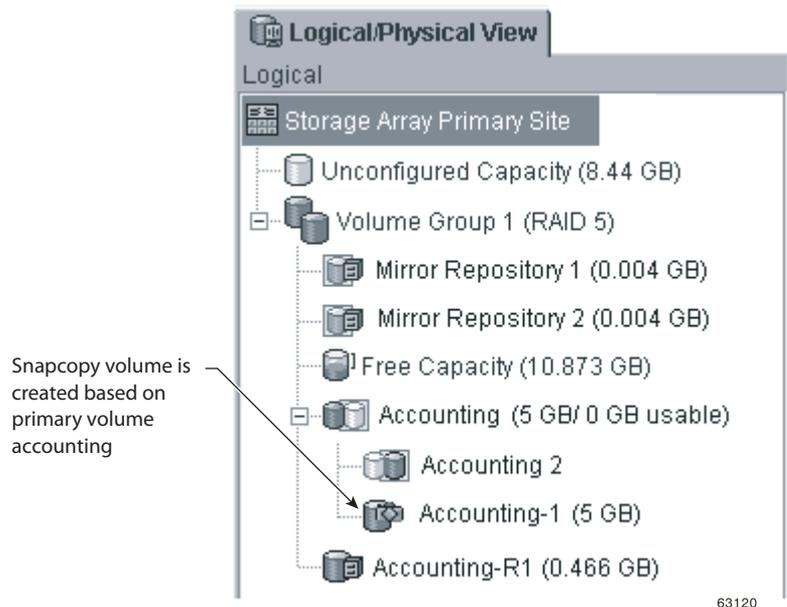


Figure 3-10 Snapshot Volume Associated with Primary Volume Accounting

If a role reversal is performed, Primary Volume Accounting is demoted to a secondary role and the snapshot volume fails, as shown in Figure 3-11. This is because the Secondary Volume Accounting 2 is now considered as the base volume of the snapshot volume, which is not permitted.

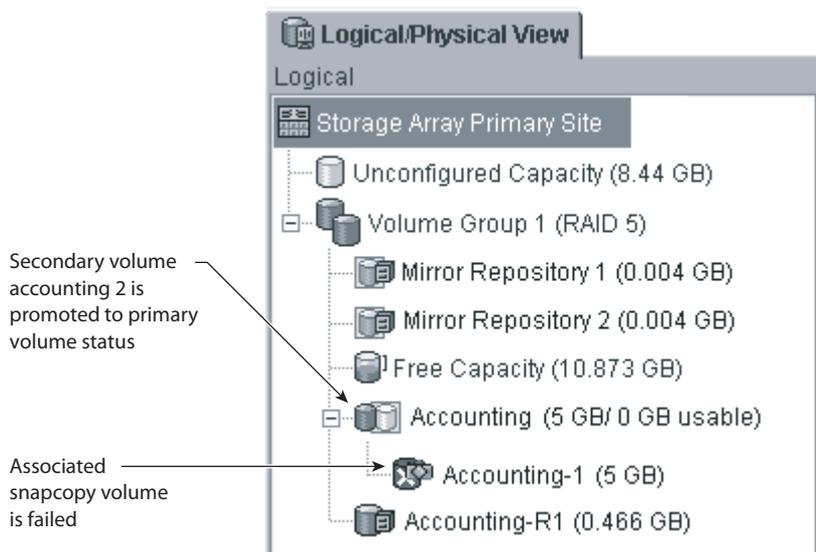


Figure 3-11 Snapshot Volume Failure Because of Role Reversal

If a snapshot volume fails as the result of a role reversal, you will need to perform another role reversal to promote the secondary volume back to a primary role. After the role reversal is complete, the snapshot volume will return to an Optimal status. For more information, see “Troubleshooting” on page 167.

Volume Copy

The Volume Copy feature copies data from one volume (the source volume) to another volume (the target volume) in a single storage system. The Volume Copy feature performs these operations:

- Copies data from drive groups that use smaller-capacity drives to drive groups that use larger-capacity drives
- Backs up data
- Restores snapshot volume data to the base volume. This feature includes a Create Copy Wizard to help you create a volume copy. This feature also includes Copy Manager to monitor volume copies after they have been created.

A primary volume in a mirror relationship can be a source volume or a target volume in a volume copy. Volume copies are not allowed on secondary volumes. If a copy of a secondary volume is required, you must perform a role reversal to change the secondary volume to a primary volume. If a role reversal is initiated during a copy in progress, the copy fails and cannot be restarted.

Note: Another method to produce a copy of the secondary volume is to create a snapshot volume of the secondary volume, and then perform a volume copy operation on the snapshot volume. This method removes the need to perform a role reversal on the secondary volume.

Dynamic Volume Expansion

Dynamic Volume Expansion (DVE) increases the capacity of a standard volume or a snapshot repository volume. The increase in capacity is achieved by using the free capacity that is available on the drive group of the standard volume or the snapshot repository volume.

This modification operation is considered to be “dynamic” because you have the ability to continually access data on drive groups, volumes, and drives throughout the entire operation.

You can perform a DVE operation on a primary volume or a secondary volume of a mirror relationship. Although the storage management software indicates that the

unsynchronized volume has increased in capacity, the usable capacity of the volume becomes the smaller of the two volumes.

Note: The mirror relationship must have an Optimal status to perform a DVE operation.

A DVE operation is not permitted on a mirror repository volume. If the mirror repository volume exists from a previous release of the storage management software and must be upgraded.

Performing Defragment Operations

Using an operating system-specific defragment utility to perform a defragment operation on a base volume with an associated snapshot volume will cause a copy-on-write of every data block in the base volume. This can cause the snapshot repository to fill before the defragment operation is completed. As a result, the snapshot volume will fail or host writes will fail to the base volume, depending on the Snapshot Volume Full condition that was originally set for the snapshot volume.

To prevent this from occurring, ensure that the snapshot repository volume's capacity is set to at least 105% of the size of the base volume before using a defragment utility. This is the minimum size needed to support a copy-on-write of every data block in the base volume, which will occur as a result of a defragment operation.

Hardware and Software Requirements

This section describes the minimum hardware requirements for the Remote Volume Mirroring feature. A list of the certified host bus adapters for each operating system is provided.

The minimum software requirements for the Remote Volume Mirroring feature are provided, together with the host operating systems certified for use with this version of the storage management software.

Hardware Requirements

The following are the minimum hardware requirements needed when configuring your system to use the Remote Volume Mirroring premium feature. The list includes components that are supplement.

- Minimum of two multi-ported storage arrays
- RAID controllers with two or more host ports
- Switches – These devices are used to connect multiple hosts with multiple storage arrays in a controlled environment. Multiple switches can be connected over a distance up to 10 km (6.25 miles) to create a fabric environment.
- Short or Medium Range Fibre Channel Cables – To connect the hosts and storage arrays to the switches.
- Long Range Fibre Channel Cables (Campus Configurations only) – These cables are up to 10 km (6.25 miles) in length. These are used to connect Fibre Channel switches to complete long distance fabrics.

Software Requirements

The following are the minimum software requirements needed when configuring your system to use the Remote Volume Mirroring premium feature.

- Total Performance Storage System Manager for Version 8.40 or later
- Controller Firmware Version 6.40 or later
- Remote Volume Mirroring Feature Key

Recommended Configurations

This section describes the key configurations supported for Remote Volume Mirroring in this version of the storage management software. The level of redundancy will be determined by the type of configuration chosen.

Highest Availability Campus Configuration

Note: The highest availability campus configuration is the recommended configuration for the Remote Mirroring feature.

This configuration has two Fibre Channel switches at both the primary site and the secondary site (four switches total), which provide for complete failover and redundancy (refer to Figure 3-12). Failures could involve Fibre Channel switches, Fibre Channel cables, and any host or storage system. Two Fibre Channel switches at each site also provide redundancy to the local site in addition to a fully redundant remote configuration. No single point of failure exists in the hardware components.

Figure 3-12 shows the highest availability campus configuration using IS4500 controller enclosures, which are four-port systems. In this configuration, use the A4 connection and the B4 connection for remote mirroring traffic instead of the A2 connection and the B2 connection. You can use controller ports A2, A3, B2, and B3 for additional host access if needed.

For detailed instructions on setting up this configuration, see “Hardware and Software Installation” on page 108.

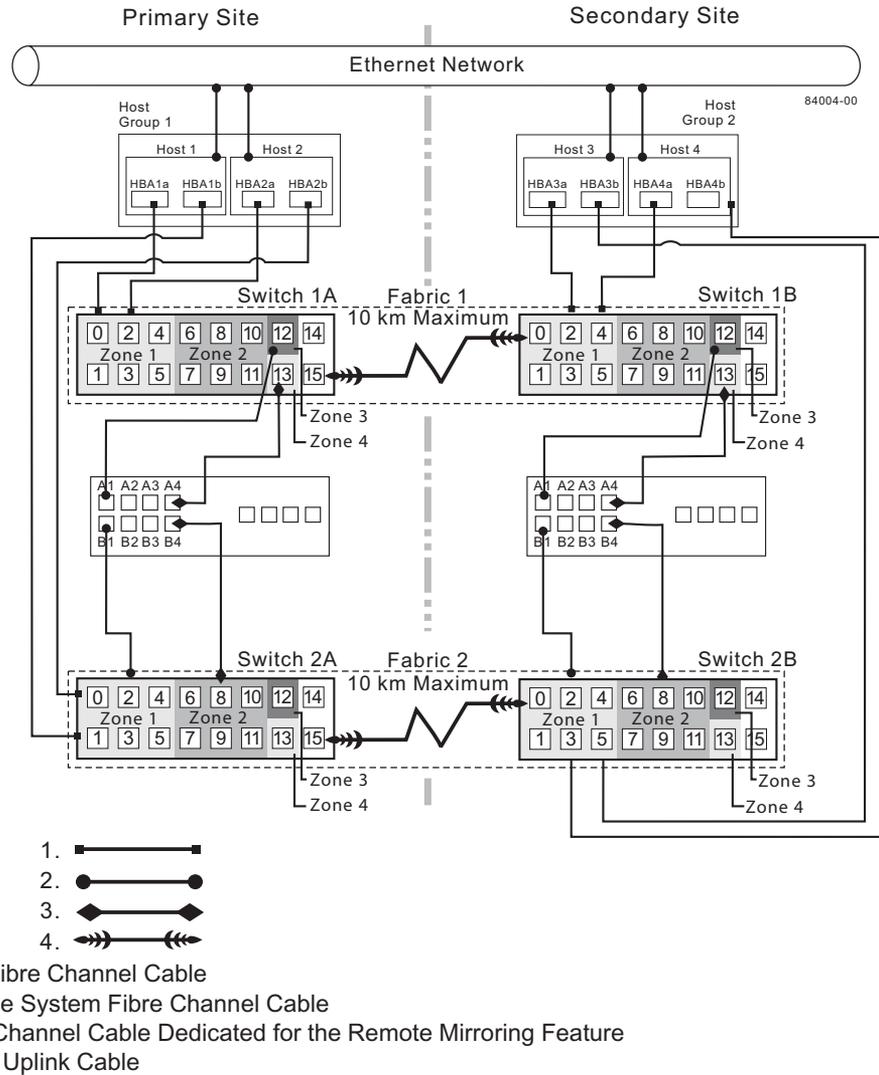


Figure 3-12 Highest Availability Campus Configuration

Campus Configuration

The Campus Configuration is a lower-cost configuration and consists of two storage arrays and two switches, connected using a Fibre Channel fabric as shown in Figure 3-13. The primary storage array and secondary storage array may have a maximum connection distance of up to 10 km (or 6.25 miles).

For detailed instructions on setting up this configuration, see “Hardware and Software Installation” on page 108.

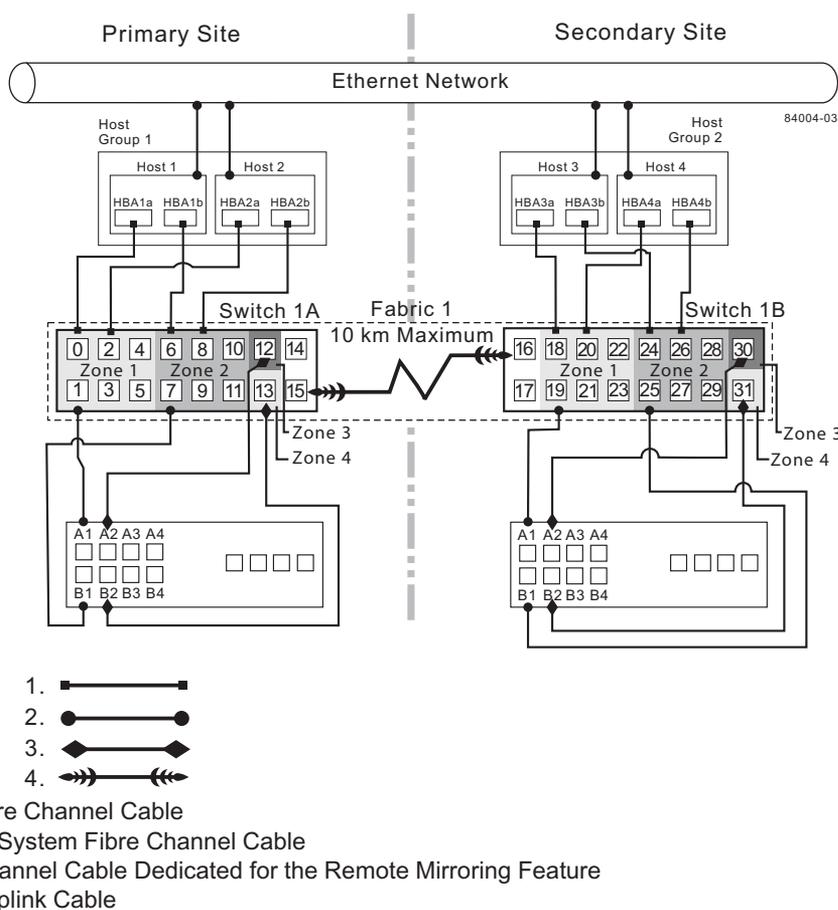


Figure 3-13 Campus Configuration

Intra-Site Configuration

The Intra-Site Configuration is the lowest-cost configuration which consists of two storage arrays and two switches, connected using a Fibre Channel fabric as shown in Figure 3-14. The primary storage array and secondary storage array may have a maximum connection distance of up to 500 m (or 0.32 miles) and would typically be located in the same building or in the same room.

For detailed instructions on setting up this configuration, see “Hardware and Software Installation” on page 108.

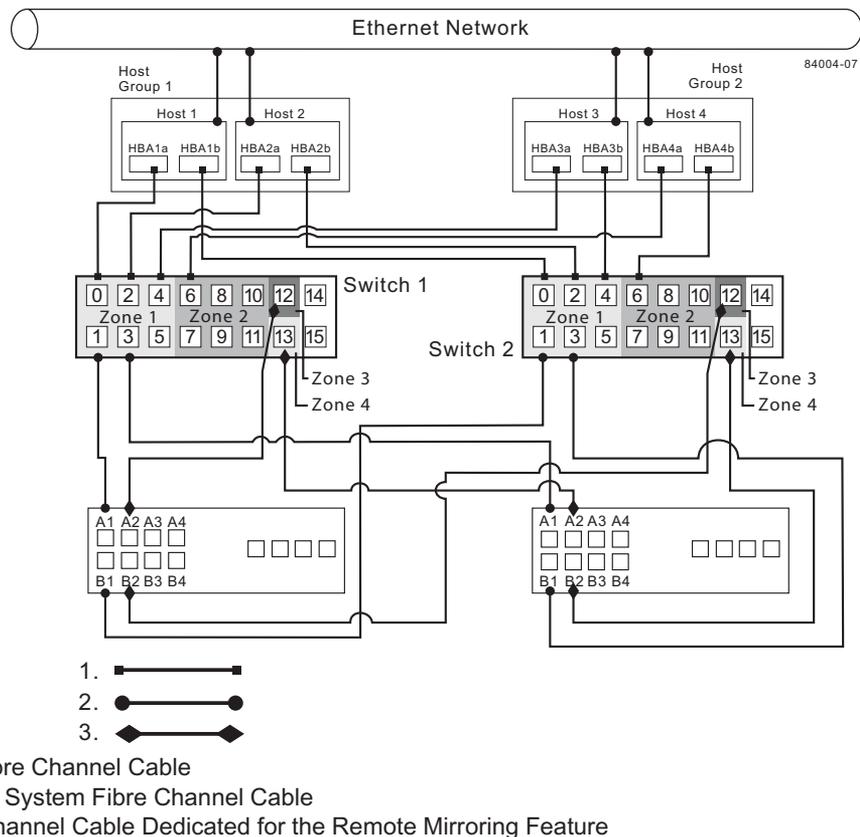


Figure 3-14 Intra-Site Configuration

Hardware and Software Checklist

Complete the following checklists to ensure that all the required hardware and software components have been configured properly.

Hardware Checklist

Complete the tasks in Table 3-1 before starting the storage management software.

Table 3-1 Hardware Checklist

Task Description	Validation Activity
1. Ensure that the minimum hardware requirements have been met. For a complete list of the hardware requirements, see “Hardware Requirements” on page 101.	<input type="checkbox"/> Completed.
2. Ensure your primary and secondary storage arrays have been properly configured.	<input type="checkbox"/> Completed.
3. Ensure that your switches and cables have been properly configured. For more information on configuring switches for use with the Remote Volume Mirroring feature, see “Hardware and Software Installation” on page 108.	<input type="checkbox"/> Completed.

Software Checklist

Complete the tasks in Table 3-2 before establishing mirror relationships in your storage arrays.

Table 3-2 Software Checklist

Task Description	Validation Activity
1. Ensure that data on the primary and secondary storage arrays participating in Remote Volume Mirroring have been backed up.	__ Completed.
2. Ensure that the correct version of firmware and storage management software has been installed.	__ Completed.
For more information on upgrading firmware or installing the storage management software, see “Hardware and Software Installation” on page 108.	
3. Ensure that the Remote Volume Mirroring premium feature has been enabled on both the primary and secondary storage arrays.	__ Completed.
For more information on enabling premium features, see “Hardware and Software Installation” on page 108 or the Array Management window online help.	
4. Ensure that the Remote Volume Mirroring feature has been activated and that a mirror repository volume has been created for each controller on the primary storage array.	__ Completed.
For more information on activating the Remote Volume Mirroring feature, see “Using the Remote Volume Mirroring Premium Feature” on page 134.	
5. Ensure that the required primary and secondary volumes have been created on the primary and secondary storage arrays.	__ Completed.
For more information on establishing a mirror relationship, see “Using the Remote Volume Mirroring Premium Feature” on page 134.	

Hardware and Software Installation

This section presents information necessary to install hardware and software to run the Remote Volume Mirroring premium feature. The beginning of the section explores the steps that need to be completed before any hardware or software installation. Next, hardware installation procedures are presented, followed by software installation procedures. This section is to be followed sequentially, from start to end.

- “Pre-Installation Requirements” on page 108
- “Hardware Installation” on page 111
- “Software Installation” on page 131

Pre-Installation Requirements

This section provides necessary pre-installation information, such as site preparation, hardware and software requirements, and concepts that should be reviewed and completed before beginning any hardware or software installation procedures.

Site Preparation

The Remote Volume Mirroring feature utilizes Fibre Channel switches creating a fabric environment for data replication. These switches require only minimal additional site preparation requirements beyond basic storage array operation.

For additional site preparation considerations concerning the Fibre Channel switches, including power requirements and physical dimensions/requirements, see the documentation provided by the switch manufacturer.

Hardware and Software Requirements

Ensure that all of the required hardware and software are available before proceeding. For a list of requirements, see “Hardware and Software Installation” on page 108.

Switch Zoning Overview

Because of possible restrictions at the host level, the supported Remote Volume Mirroring configurations contain Fibre Channel switches zoned so a single host adapter can access only one controller per storage array. Additionally, all configurations utilize a separate zone for the ports reserved for the Remote Volume Mirroring feature.

Note: Do not zone the uplink port (E_port) that connects (cascades) switches within a fabric.

Switch zoning configurations are typically set up using management software provided by the manufacturer of the Fibre Channel switch. This software should be included with the materials provided when the switch was purchased.

When two or more Fibre Channel switches are cascaded together, the switch management software will combine the ports for all switches that are linked. For example, if two 16-port Fibre Channel switches are cascaded with a physical connection using a Fibre Channel cable, the switch management software will show ports 0 through 31 participating in the fabric rather than 2 switches each with ports 0 through 15. Therefore, a zone that is created containing any of these ports can exist on multiple cascaded switches.

In Figure 3-15, the top set of switches are on the same network and cascaded, therefore Zone 1 is the same zone on Switch 1A as Zone 1 on Switch 1B. The bottom set of switches are on the same network, but are not cascaded. Even though both contain a Zone 1 (shown as Zone A in Switch 2), these zones are independent of each other.

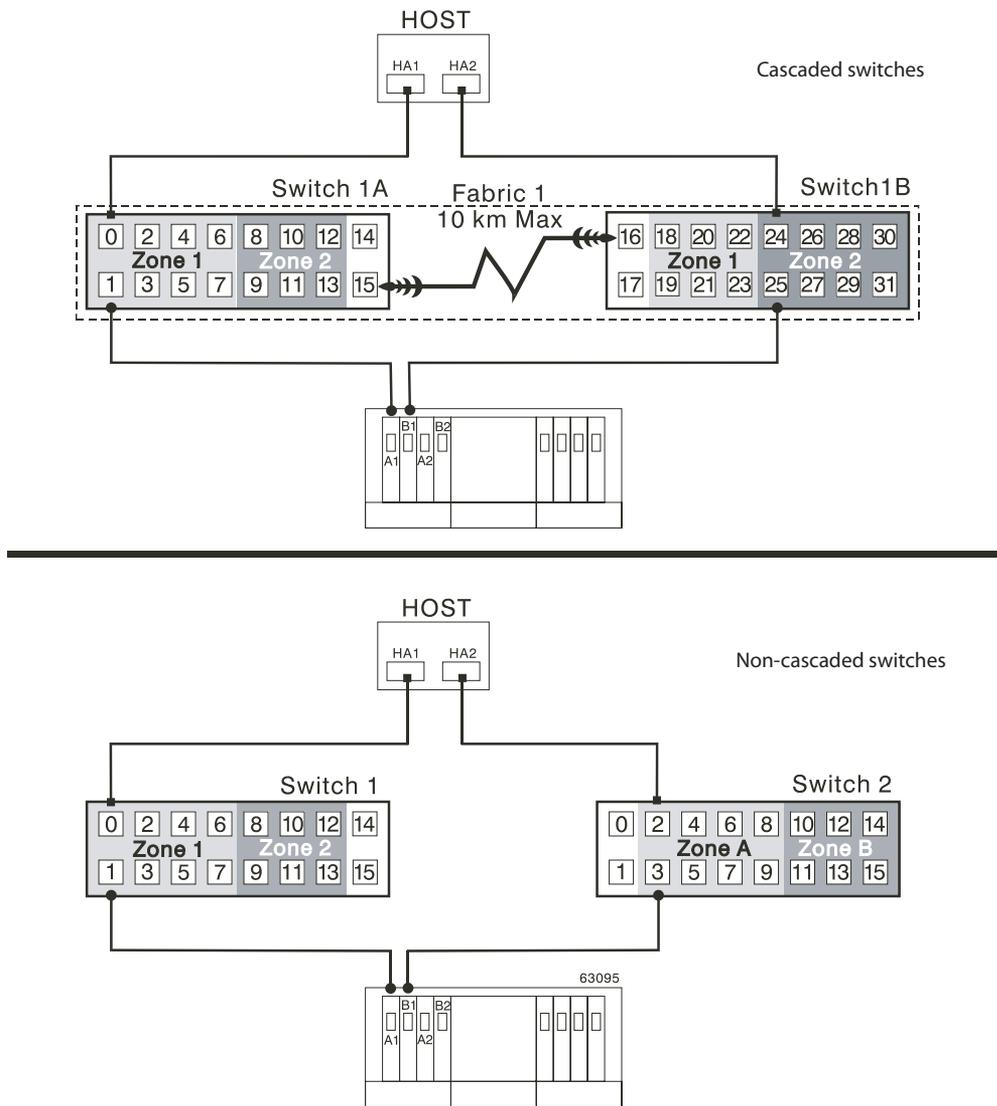


Figure 3-15 Switch Zoning in Cascaded and Non-Cascaded Fibre Channel Switches

For more information about Fibre Channel switch zoning or setting up a zone configuration, see the manufacturer's documentation provided with the switch.

Due to the varying Remote Volume Mirroring configurations, the switch zone settings are presented preceding each configuration illustration in this chapter.

Hardware Installation

This section describes the procedures necessary to properly connect and configure one or more storage arrays for use with the Remote Volume Mirroring feature.

The configurations covered in this section are as follows.

- **Highest Availability Campus Configuration** – This is the recommended configuration for the Remote Volume Mirroring feature. Four total switches (two at each site) provide for complete redundancy in the Fibre Channel switches and fabrics in addition to all storage array components and hosts. There is no single point of failure in this configuration. To begin installing hardware for this configuration, go to the next section, “Highest Availability Campus Configuration” on page 101.
- **Campus Configuration** – This configuration is not fully redundant because the number of Fibre Channel switches is reduced from four to two and the number of fabrics from two to one. Consequently, this makes the switches and fabrics a single point of failure. A switch failure does not usually result in a loss of data, but does affect data synchronization until the error is corrected. Therefore, the Highest Availability Campus Configuration is the recommended configuration because data synchronization can continue for any single switch failure. This configuration essentially allows the minimum required components to successfully operate the Remote Volume Mirroring premium feature between two sites up to 10 km (6.25 miles). To begin installing hardware for this configuration, go to “Campus Configuration” on page 104.
- **Intra-Site Configuration** – This configuration is similar to the Campus Configuration because there are only two switches. However, no multiple-switch fabrics exist in this configuration. Since all hosts and storage arrays are in proximity of both Fibre Channel switches, full redundancy can be obtained within the site since each controller path is connected to a different switch. Despite being a fully redundant configuration, the Highest Availability Campus Configuration is still recommended because of disaster recovery implications. The single site creates a

single point of failure if a disaster affects the entire site. To begin installing hardware for this configuration, go to “Intra-Site Configuration” on page 105.

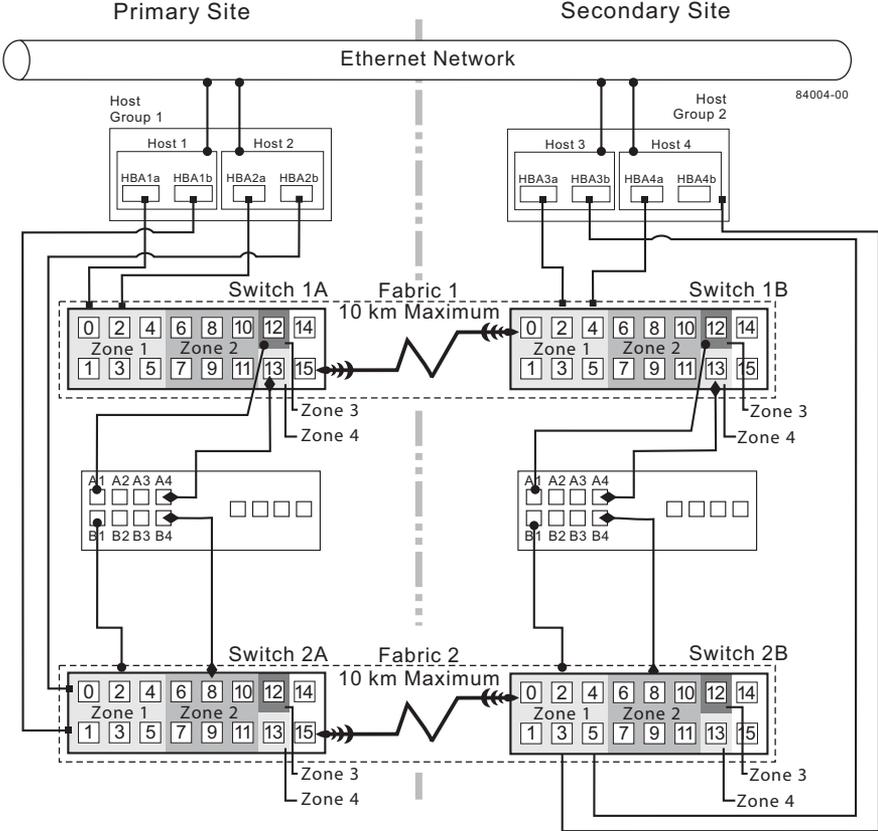
For a list of required materials to complete any hardware installation, see “Hardware and Software Installation” on page 108.

Highest Availability Campus Configuration

Note: The Highest Availability Campus Configuration is the recommended configuration for the Remote Volume Mirroring feature.

This configuration has two Fibre Channel switches at both the primary and secondary sites to provide for complete failover and redundancy for failures in Fibre Channel switches and cables, in addition to any host or storage array failure. Two Fibre Channel switches at each site also provide redundancy to the local site in addition to a fully redundant remote configuration. There is no single point of failure in the hardware components.

Figure 3-16 shows the Highest Availability Campus Configuration after installation is complete.



- 1. —————
 - 2. —●—●—
 - 3. —◆—◆—
 - 4. —>>—<<—
1. Host Fibre Channel Cable
 2. Storage System Fibre Channel Cable
 3. Fibre Channel Cable Dedicated for the Remote Mirroring Feature
 4. Fabric Uplink Cable

Figure 3-16 Highest Availability Campus Configuration

Zoning for the Highest Availability Campus Configuration

This configuration allows for a separate zone for each reserved port for the Remote Volume Mirroring feature.

The switches do not need to be zoned exactly as presented in this configuration. However, the following are requirements when zoning switches for the Highest Availability Campus Configuration.

- There s a total of four zones in this configuration.
 - Zones 1 and 3 will exist on Fabric 1 (Switch 1A at the primary Site, and Switch 1B at the secondary site).
 - Zones 2 and 4 will exist on Fabric 2 (Switch 2A at the primary Site, and Switch 2B at the secondary site).
- Configure the zones on the switch so that there is one port per zone for a storage array connection and one port per zone for each host.
- Switches are zoned so that a single host adapter can access only one controller per storage array.

Note: Do not zone the uplink ports (E_ports) on any of the Fibre Channel switches.

The switches in Figure 3-17 contain 16 ports each, which leaves many unused ports per switch. The remaining ports can be distributed among the other zones. However, it is recommended that most remaining ports be assigned to the zones containing the host connections (Zones 1 and 2 in Figure 3-17). This will allow easy set up for additional hosts to connect to the environment.

Figure 3-17 shows how the four switches are zoned for the Highest Availability Campus Configuration.

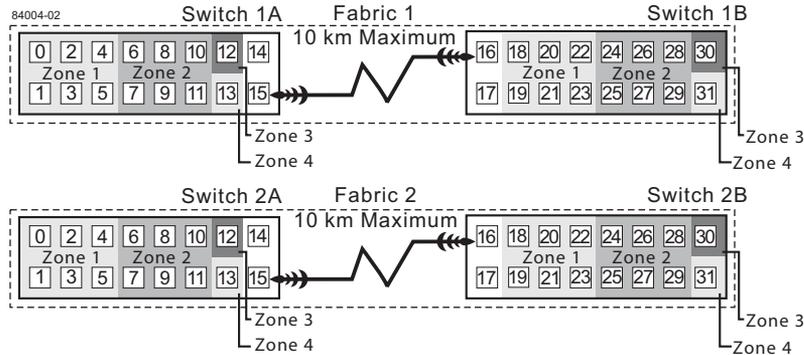


Figure 3-17 Switch Zoning for the Highest Availability Campus Configuration

Given the requirements and the zoning shown in Figure 3-17, ensure that all four switches are properly zoned before proceeding. For more information on Zoning switches, see “Switch Zoning Overview” on page 109.

Cabling for the Highest Availability Campus Configuration

Note: Begin the installation at the primary site. Repeat these steps for the secondary site when instructed to do so.

After the four Fibre Channel Switches are properly zoned, complete the following procedure to set up the Highest Availability Campus Configuration for the Remote Volume Mirroring feature.

Note: All connections are completed using Fibre Channel cables of appropriate length.

1. Is Remote Volume Mirroring feature equipment being added to an existing storage array environment?
 - Yes – Stop I/O from all hosts before proceeding. Go to 2.
 - No – This is a new storage array installation. Go to 3.

Caution: Possible hardware damage – DC-powered controller enclosures, controller-drive enclosures, and drive enclosures have special power procedures that you must follow beyond what is usually expected. Refer to the user guide for your enclosure to get specific poweroff and power-on procedures.

2. Power down all storage arrays, hosts, Fibre Channel switches, and all other equipment in the storage array environment.
3. Ensure that cabling between all controller enclosures and drive enclosures is complete.

Note: Depending on which site is being configured, Switch 1 represents Switch 1A for the primary site, and Switch 1B for the secondary site. This applies to Switch 2 as well.

4. Connect the primary host bus adapter (Hx1) for each local host to an available port in Zone 1 of Switch 1 (Figure 3-18).

Note: The cables can be connected to any port in the proper zone of the switch.

5. Connect the secondary host bus adapter (Hx2) for each host at this site to an available port in Zone 2 of Switch 2 (Figure 3-18).

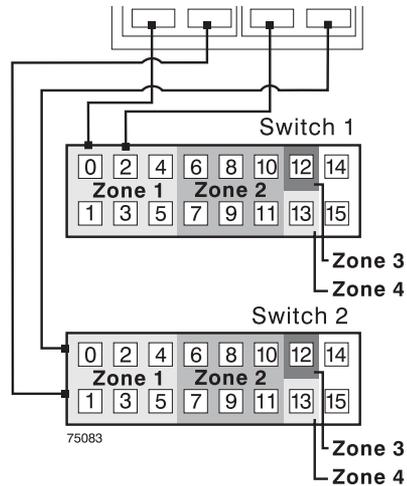


Figure 3-18 Host Bus Adapter Connections to Fibre Channel Switches

6. Connect controller port A1 of the storage array to an available port in Zone 1 of Switch 1 (Figure 3-19).
7. Connect controller port B1 of the storage array to an available port in Zone 2 of Switch 2 (Figure 3-19).
8. Connect controller port A2 of the storage array to an available port in Zone 3 of Switch 1 (Figure 3-19).
9. Connect controller port B2 of the storage array to an available port in Zone 4 of Switch 2 (Figure 3-19).

Note: Controller ports A2 and B2 are reserved for mirror relationship synchronization upon activation of the Remote Volume Mirroring feature. For more information, see “Connectivity and Input/Output” on page 91.

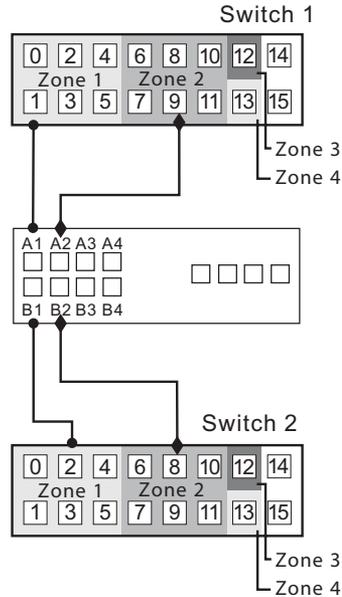


Figure 3-19 Storage Array Connections to Fibre Channel Switches

Note: Controller port A2 and controller port B2 are reserved for mirror relationship synchronization upon activation of the Remote Mirroring feature. In a four-host-port system, controller port A4 and controller port B4 are reserved. For more information see “Connectivity and Input/Output” on page 91.

10. The primary site cabling is now complete. Is the secondary site cabling complete?
 - No – Repeat steps 1 through 9 for the secondary site.
 - Yes – Go to step 11.
11. Complete the fabric environment for Switch 1 by connecting Switch 1A with Switch 1B. Connect each end of a long range Fibre Channel cable (up to 10 km in length) to an unzoned port in each switch (Figure 3-20).
12. Repeat step 11 for Switch 2A and 2B to complete the fabric environment for Switch 2 (Figure 3-20).

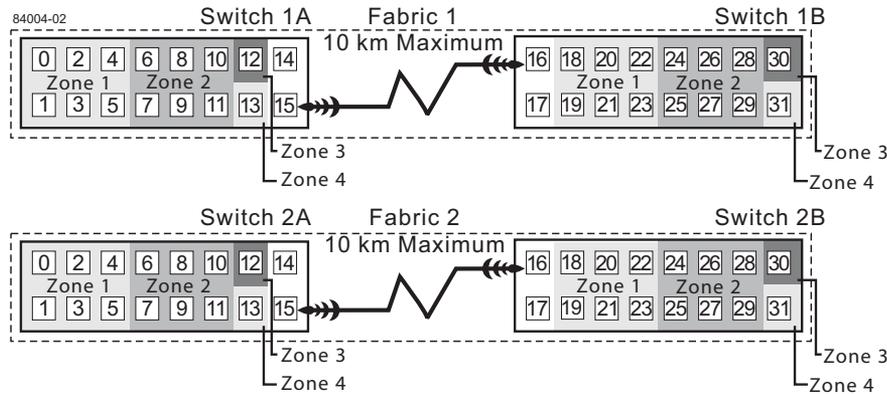


Figure 3-20 Connecting Switches over a Remote Network to Complete Fabric Environments

13. Cabling for the Highest Availability Campus Configuration is complete. Repeat 4 through 10 for any additional storage arrays that exist in the same cabinet that will use the Remote Volume Mirroring feature.

Caution: Possible hardware damage – DC-powered controller enclosures, controller-drive enclosures, and drive enclosures have special power procedures that you must follow beyond what is usually expected. Refer to the user guide for your enclosure to get specific poweroff and power-on procedures.

14. Power up all storage arrays, hosts, Fibre Channel switches, and any other hardware at both sites that was powered down in 2.
15. Hardware installation is complete. The next step will be to configure the storage management software to support mirror relationships. Go to “Software Installation” on page 131.

Campus Configuration

This configuration offers the same functionality as the Highest Availability Campus Configuration, but contains only one switch at each site, rather than two. The configuration is still redundant for host bus adapters, controllers, and Remote Volume Mirroring ports, but is a single point of failure for switches. If a switch at either site fails, the Remote Volume Mirroring Feature cannot function. For this reason, the Highest Availability Campus Configuration is highly recommended for total environment redundancy.

Figure 3-21 shows the Campus Configuration after installation is complete.

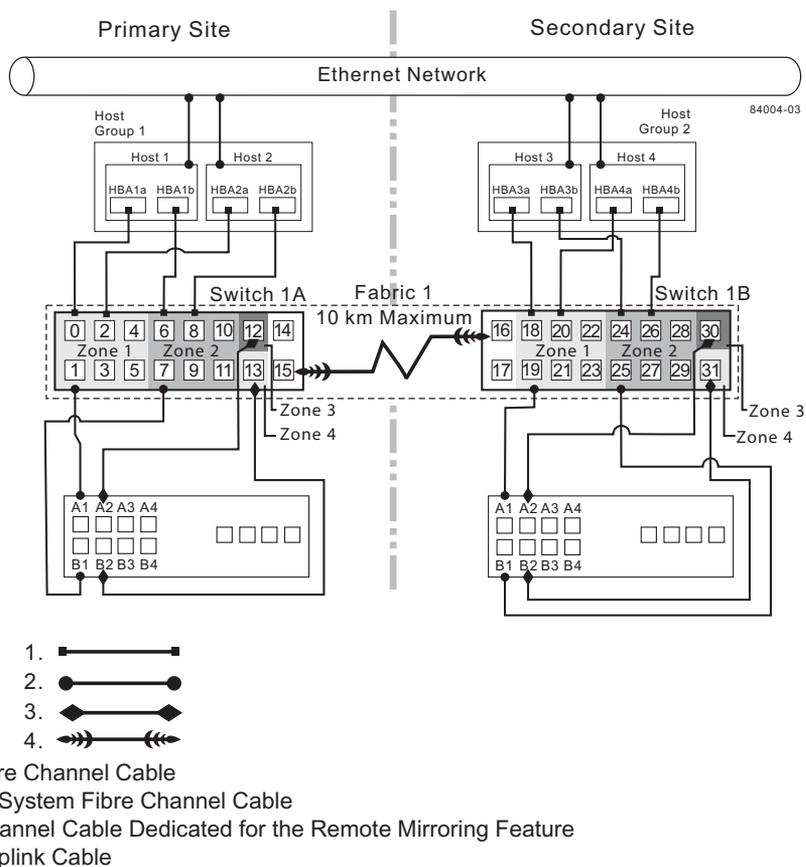


Figure 3-21 Campus Configurations

Switch Zoning for the Campus Configuration

This configuration allows for a separate zone for each reserved port for the Remote Volume Mirroring feature.

The switches do not need to be zoned exactly as presented in this configuration. However, the following are requirements when zoning switches for the Campus Configuration.

- There will be a total of four zones in this configuration.
 - All zones will exist on Fabric 1 (Switch 1A at the primary site, and Switch 1B at the secondary site).
 - Zones 3 and 4 will be reserved for the dedicated Remote Volume Mirroring Feature connections.
- Configure the zones on the switches so that there is one port per zone for a storage array connection and one port per zone for each host.

Important: Do not zone the uplink ports (E_ports) on any of the Fibre Channel switches.

- Switches are zoned so that a single host adapter can access only one controller per storage array.

The switches in Figure 3-22 contain 16 ports each, which leaves many unused ports per switch. The remaining ports can be distributed among the other zones. However, it is recommended that most remaining ports be assigned to the zones containing the host connections (Zone 1 in Figure 3-22). This will allow easy set up for additional hosts to connect to the environment.

Figure 3-22 shows how the two switches are zoned for the Campus Configuration.

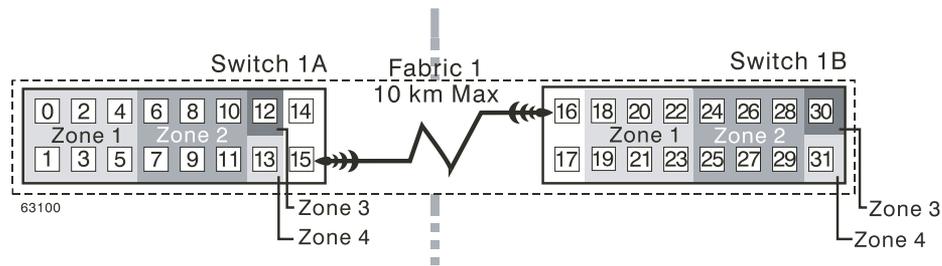


Figure 3-22 Switch Zoning for the Campus Configuration

Given the requirements and the zoning example shown in Figure 3-22, ensure that both switches are properly zoned before proceeding. For more information on Zoning switches, see “Switch Zoning Overview” on page 109.

Cabling for the Campus Configuration

Note: Begin the installation at the primary site. Repeat these steps for the secondary site when instructed to do so.

After both Fibre Channel Switches are properly zoned, complete the following procedure to set up the Campus Configuration for the Remote Volume Mirroring feature.

Note: All connections are completed using Fibre Channel cables of appropriate length.

1. Is Remote Volume Mirroring feature equipment being added to an existing storage array environment?
 - **Yes** – Stop I/O from all hosts before proceeding. Go to step 2.
 - **No** – This is a new storage array installation. Go to step 3.

Caution: Possible hardware damage – DC-powered controller enclosures, controller-drive enclosures, and drive enclosures have special power procedures that you must follow beyond what is usually expected. Refer to the user guide for your enclosure to get specific poweroff and power-on procedures.

2. Power down all storage arrays, hosts, Fibre Channel switches, and all other equipment in the storage array environment.
3. Ensure that basic cabling between all controller enclosures and drive enclosures is complete.

Note: Depending on which site is being configured, Switch 1 represents Switch 1A for the primary site, and Switch 1B for the secondary site.

4. Connect the primary host bus adapter (Hx1) for each host at this site to an available port in Zone 1 of Switch 1 (Figure 3-23).

Note: The cables can be connected to **any** port in the proper zone of the switch.

5. Connect the secondary host bus adapter (Hx2) for each host at this site to an available port in Zone 2 of Switch 1 (Figure 3-23).

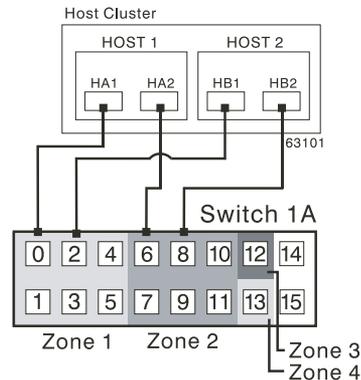


Figure 3-23 Host Bus Adapter Connections to Fibre Channel Switches

6. Connect controller port A1 of the storage array to an available port in Zone 1 of Switch 1 (Figure 3-24).
7. Connect controller port B1 of the storage array to an available port in Zone 2 of Switch 1 (Figure 3-24).
8. Connect controller port A2 of the storage array to an available port in Zone 3 of Switch 1 (Figure 3-24).

9. Connect controller port B2 of the storage array to an available port in Zone 4 of Switch 1 (Figure 3-24).

Note: Controller ports A2 and B2 are reserved for mirror relationship synchronization upon activation of the Remote Volume Mirroring feature. For more information, see “Connectivity and Input/Output” on page 91.

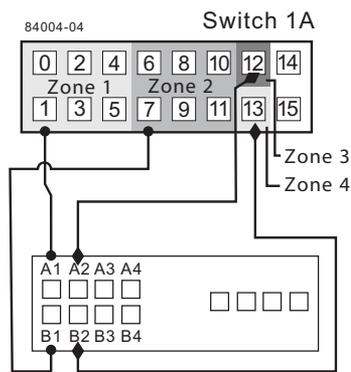


Figure 3-24 Storage Array Connections to Fibre Channel Switches

10. The primary site cabling is now complete. Is the secondary site cabling complete?
 - **No** – Repeat steps 1 through 9 for the secondary site.
 - **Yes** – Go to step 11.
11. Complete Fabric 1 by connecting Switch 1A with Switch 1B. Connect each end of a long range Fibre Channel cable (up to 10 km in length) to an unzoned port in each switch (Figure 3-25).

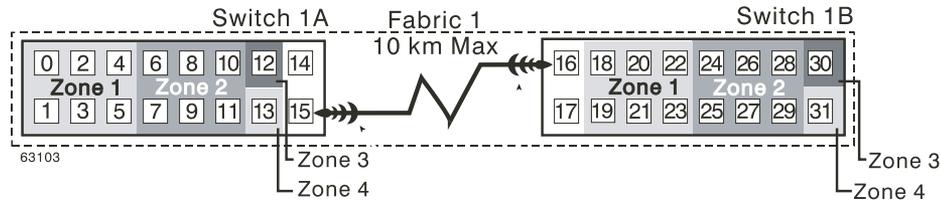


Figure 3-25 Connecting Switches over a Remote Network to Complete the Fabric Environment

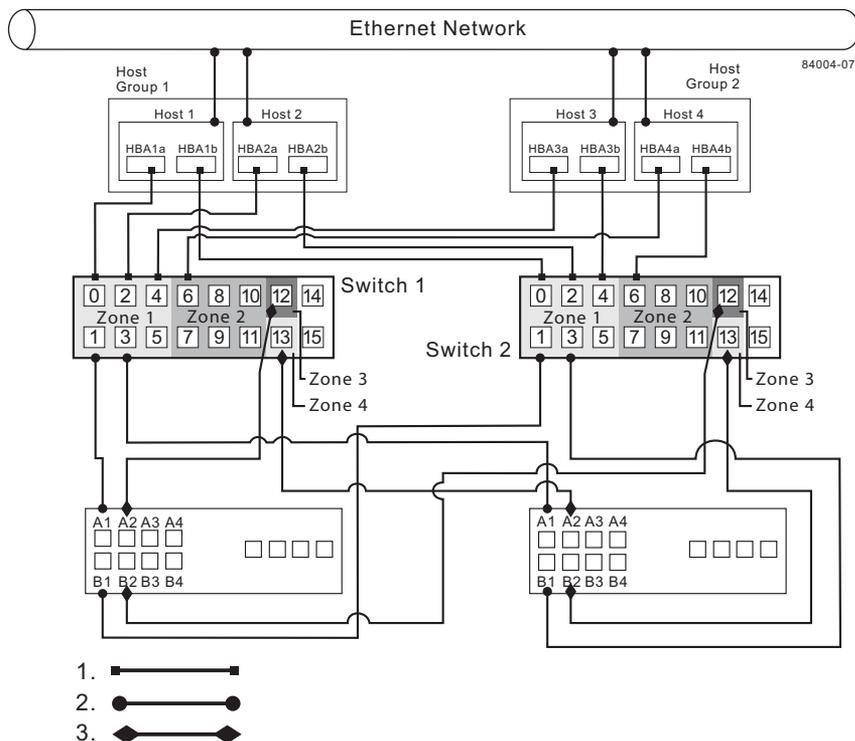
12. Cabling for the Campus Configuration is complete. Repeat steps 4 through 10 for any additional storage arrays that exist in the same cabinet that will use the Remote Volume Mirroring feature.
13. Power up all storage arrays, hosts, Fibre Channel switches, and any other hardware at both sites that was powered down in step 2.
14. Hardware installation is complete. The next step will be to configure the storage management software to support mirror relationships. Go to “Software Installation” on page 131.

Intra-Site Configuration

This configuration is used in environments where a long distance fabric is not required because of the close proximity of the hosts and storage arrays. The configuration is still redundant for host bus adapters, controllers, Remote Volume Mirroring ports, and switches, but is a single point of failure for the site since all equipment can be destroyed by the same disaster. For this reason, the Highest Availability Campus Configuration is highly recommended for total environment redundancy.

Note: A switch failure in this configuration does not affect data access, however, an unsynchronized mirror state may occur as a result.

Figure 3-26 shows the Intra-Site Configuration after installation is complete.



1. Host Fibre Channel Cable
2. Storage System Fibre Channel Cable
3. Fibre Channel Cable Dedicated for the Remote Mirroring Feature

Figure 3-26 Intra-Site Configuration

Switch Zoning for the Intra-Site Configuration

This configuration is designed for switch redundancy. However, the switches are not cascaded, and are therefore independent of each other. The switches do not need to be zoned exactly as presented in this configuration. However, the following are requirements when zoning switches for the Intra-Site Configuration.

Note: Switch 2 contains Zones A and B to show that there is no relationship between Switch 1 and Switch 2, contrary to the other configurations.

- There will be a total of four zones in this configuration.
 - Zones 1 and 2 will exist on Switch 1.
 - Zones A and B will exist on Switch 2.
- Configure the zones on the switch so that there is one port per zone for a storage array connection and one port per zone for each host.
- Switches are zoned so that a single host adapter can only access one controller per storage array.

The switches in Figure 3-27 contain 16 ports each, which leaves many unused ports per switch. The remaining ports can be distributed among the other zones. However, it is recommended that most remaining ports be assigned to the zones containing the host connections (Zone 1 and Zone A in Figure 3-27). This will allow easy set up for additional hosts to connect to the environment.

In this example, the switches use half of the ports for each zone for simplicity, although Zone 2 and Zone B require less ports.

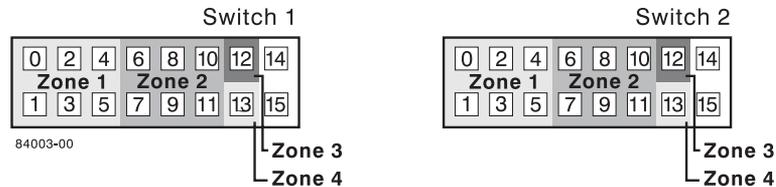


Figure 3-27 Switch Zoning for the Intra-Site Configuration

Given the requirements and the zoning example shown in Figure 3-27, ensure both switches are properly zoned before proceeding. For more information on zoning switches, see “Switch Zoning Overview” on page 109.

Cabling for the Intra-Site Configuration

After both Fibre Channel Switches are properly zoned, complete the following procedure to set up the Intra-Site Configuration for the Remote Volume Mirroring feature.

Note: All connections are completed using Fibre Channel cables of appropriate length.

1. Is Remote Volume Mirroring feature equipment being added to an existing storage array environment?
 - **Yes** – Stop I/O from all hosts before proceeding. Go to step 2.
 - **No** – This is a new storage array installation. Go to step 3.

Caution: Possible hardware damage – DC-powered controller enclosures, controller-drive enclosures, and drive enclosures have special power procedures that you must follow beyond what is usually expected. Refer to the user guide for your enclosure to get specific poweroff and power-on procedures.

2. Power down all storage arrays, hosts, Fibre Channel switches, and all other equipment in the storage array environment.
3. Ensure that basic cabling between all controller enclosures and drive enclosures is complete on both storage arrays.
4. Connect the primary host bus adapter for each host (Hx1) to an available port in Zone 1 of Switch 1 (Figure 3-28).

Note: The cables can be connected to **any** port in the proper zone of the switch.

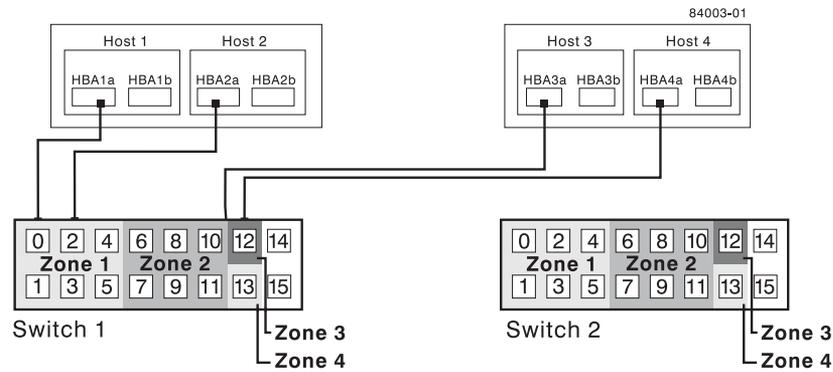


Figure 3-28 Primary Host Bus Adapter Connections to Fibre Channel Switches

5. Connect the secondary host bus adapter (Hx2) for each host to an available port in Zone A of Switch 2 (Figure 3-29).

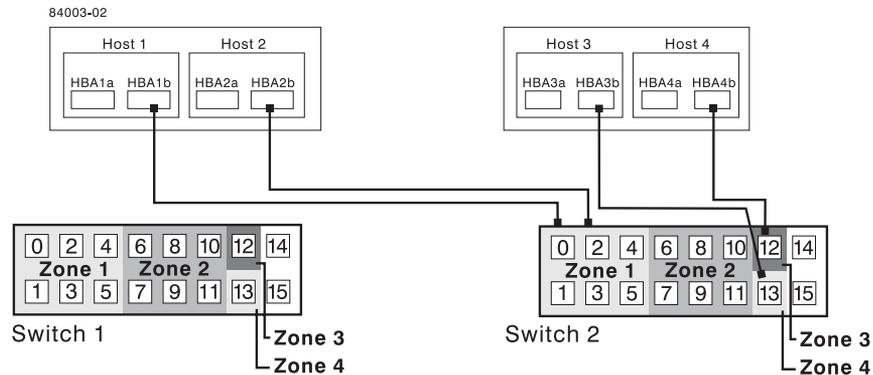


Figure 3-29 Secondary Host Bus Adapter Connections to Fibre Channel Switches

6. Connect controller port A1 of the primary storage array to an available port in Zone 1 of Switch 1 (Figure 3-30).
7. Connect controller port B1 of the primary storage array to an available port in Zone A of Switch 2 (Figure 3-30).

8. Connect controller port A2 of the primary storage array to an available port in Zone 2 of Switch 1 (Figure 3-30).
9. Connect controller port B2 of the primary storage array to an available port in Zone B of Switch 2 (Figure 3-30).

Note: Controller Ports A2 and B2 are reserved for mirror relationship synchronization upon activation of the Remote Volume Mirroring feature. For more information, see “Connectivity and Input/Output” on page 91.

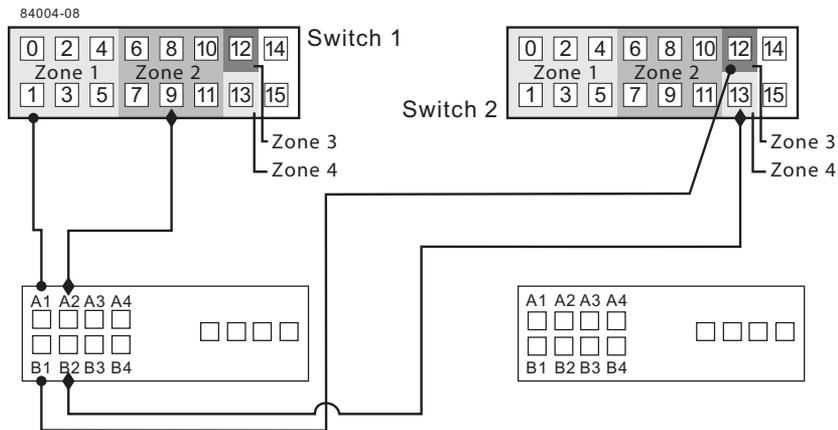


Figure 3-30 Primary Storage Array Connections to Fibre Channel Switches

10. Connect controller port A1 of the secondary storage array to an available port in Zone 1 of Switch 1 (Figure 3-31).
11. Connect controller port B1 of the secondary storage array to an available port in Zone A of Switch 2 (Figure 3-31).
12. Connect controller port A2 of the secondary storage array to an available port in Zone 2 of Switch 1 (Figure 3-31).
13. Connect controller port B2 of the secondary storage array to an available port in Zone B of Switch 2 (Figure 3-31).

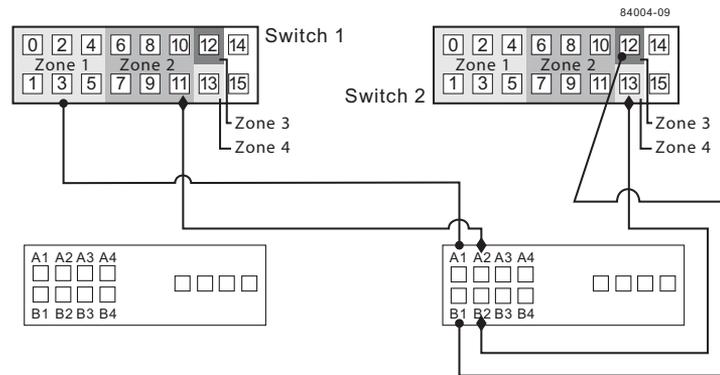


Figure 3-31 Secondary Storage Array Connections to Fibre Channel Switches

14. Cabling for the Intra-Site Configuration is complete. Repeat steps 4 through 13 for any additional storage arrays that exist in the same cabinet that will use the Remote Volume Mirroring feature.
15. Power up all storage arrays, hosts, Fibre Channel switches, and any other hardware that was powered down in step 2.
16. Hardware installation is complete. The next step will be to configure the storage management software to support mirror relationships. Go to “Software Installation” on page 131.

Software Installation

The Remote Volume Mirroring premium feature is automatically included as part of the SGI InfiniteStorage System Manager Enterprise Edition (ISSM-EE) installation. However, the feature is not enabled or activated by default after installation.

Note: All of the storage systems that participate in mirror relationships must have a supported firmware version installed. For information about the supported firmware version for this release, refer to the firmware CD that came with your storage system, or contact a Customer and Technical Support representative.

- For information about installing ISSM-EE, refer to the *SGI InfiniteStorage ISSM-EE Administration Guide* publication number, 007-5478-00x.
- For information about enabling or activating the Remote Volume Mirroring premium feature, see “Enabling and Activating the Remote Volume Mirroring Feature” on page 134.

Verifying Installation and Configuration

All hardware and software is now set up to support mirror relationships through the Remote Volume Mirroring feature. However, before creating mirror relationships, ensure that all required components are working properly.

Complete the following procedure for each storage array that will participate in any aspect of a mirror relationship to verify correct configuration.

1. Start the storage management software.

For more information about starting the storage management software, see the *SGI InfiniteStorage ISSM-EE Administration Guide* for operating system specific procedures.

2. Is this the first time the storage management software has been started?

- Yes – Go to step 3.
- No – Go to step 4.

3. Start the storage management software (for the first time).

- a. Select **OK** when prompted for Automatic Discovery of devices.
- b. Ensure that all storage arrays designated to participate in mirror relationships are displayed in the Device Tree view of the **Enterprise Management** window.

Note: If the storage arrays do not appear, see “Troubleshooting” on page 167.

4. Start the storage management software (not for the first time).

- a. From the **Enterprise Management** window, select **Tools > Automatic Discovery**.
- b. Select **OK**.
- c. Verify that all storage arrays designated to participate in mirror relationships are displayed in the Device Tree view of the **Enterprise Management** window.

- d. If the storage arrays do not appear, see “Troubleshooting” on page 167.
5. From the **Enterprise Management** window, select a storage array that will participate in mirror relationships.
6. Select **Tools > Manage Device** to open the **Array Management** window for that storage array.
7. From the **Array Management** window, select **Help > About**.
8. Verify that the version of the storage management software is 9.12 or later. If this requirement is not met, contact Technical Support for storage management software upgrade procedures.
9. Select **OK**.
10. From the **Array Management** window, select **View > Storage Array Profile**.
11. Complete the following information based on the current firmware version of the controllers:
 - Earlier than 6.12 – Upgrade to the firmware version 6.12 or later. For firmware upgrade procedures, refer to the product release notes that were provided with the TPSSM or ISSM-EE CD-ROM. When you are finished, go to step 12.
 - 6.12 or later – The firmware is currently at the supported level, so no upgrade is necessary. Go to step 12.
12. Observe the Remote Mirroring premium feature icon in the premium feature status area. The icon represents a Disabled/Deactivated status of the Remote Volume Mirroring feature.

To begin working with Mirror Relationships, go to “Using the Remote Volume Mirroring Premium Feature” on page 134.

Using the Remote Volume Mirroring Premium Feature

This section consists of procedures necessary to perform the major tasks associated with the Remote Volume Mirroring feature. The section is organized sequentially; after the feature has been enabled and activated, mirror relationships can be created and then maintained. The final subsection is provided for reference only and may be required when a mirror relationship is no longer needed.

- “Enabling and Activating the Remote Volume Mirroring Feature” on page 134
- “Creating Mirror Relationships” on page 142
- “Maintaining Mirror Relationships” on page 145
- “Deactivating and Disabling the Remote Volume Mirroring Feature” on page 156

Enabling and Activating the Remote Volume Mirroring Feature

The Remote Volume Mirroring premium feature must be enabled and activated before any mirror relationships can be created. First, determine the status of the Remote Volume Mirroring feature. After determining the status, the feature will need to be enabled and activated in order to perform all functions of the Remote Volume Mirroring feature.

Determining Remote Volume Mirroring Feature Status

There are four possible statuses of the Remote Volume Mirroring feature:

- Disabled/Deactivated
- Disabled/Activated
- Enabled/Deactivated
- Enabled/Activated.

Note: The Remote Volume Mirroring feature must be in an Enabled/Activated state on the secondary storage array as well as the primary storage array to utilize all the functionality of the feature.

The current status can be determined by hovering the mouse pointer over the Remote Volume Mirroring feature icon in the premium feature status area of the **Array Management** window, as shown in Figure 3-32.

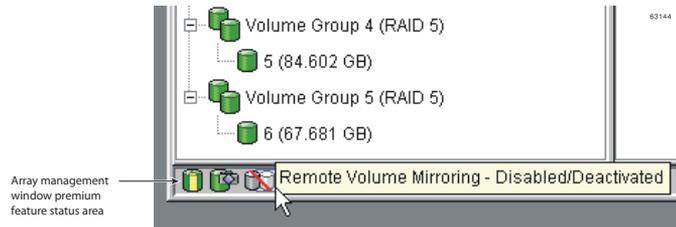


Figure 3-32 Remote Volume Mirroring Feature Status

Note: The Remote Volume Mirroring feature status for the primary storage array is managed independently from the secondary storage array. To determine the status for both storage arrays, select each storage array independently, then determine the status.

The status of the Remote Volume Mirroring feature is one of the following.

- **Disabled/Deactivated** – No functions of the Remote Volume Mirroring feature can be performed. Full Remote Volume Mirroring feature functionality is not available until the feature is enabled and activated. The icon for this state is displayed in the premium feature status area as a cylinder with a mirrored reflection. The icon has a red slash through it to show it is disabled, and is greyed out to show it is deactivated. An example of this state is illustrated in Figure 3-32. For more information about enabling the Remote Volume Mirroring feature, see “Enabling the Remote Volume Mirroring Feature” on page 136. To activate the feature, see “Activating the Remote Volume Mirroring Feature” on page 137.
- **Disabled/Activated** – The Remote Volume Mirroring feature is disabled, preventing new mirror relationships from being created. However, any pre-existing mirror relationships can be maintained with all functions of the Remote Volume Mirroring feature. The icon displayed for this state is similar to the Disabled/Deactivated state, because the red slash remains to show the feature is disabled. However, the cylinder and its reflection now appear green to show the feature is active. For more information about enabling the Remote Volume Mirroring feature, see “Enabling the Remote Volume Mirroring Feature” on page 136.

- **Enabled/Deactivated** – The Remote Volume Mirroring feature is enabled, but not activated. The icon displayed for this state is similar to the Disabled/Deactivated state, because the cylinder and its reflection remain grey. However, the red slash no longer appears, therefore, the feature is enabled. To activate the feature, see “Activating the Remote Volume Mirroring Feature” on page 137.
- **Enabled/Activated** – The Remote Volume Mirroring feature is enabled and active. Mirror relationships can be created and maintained. The icon for this state is displayed in the premium feature status area as a green cylinder with a mirrored reflection. For more information about creating a mirror relationship, see “Creating Mirror Relationships” on page 142.

Enabling the Remote Volume Mirroring Feature

If the current status of the Remote Volume Mirroring feature is Disabled/Deactivated or Disabled/Activated, complete the following to enable the Remote Volume Mirroring feature.

1. From the **Array Management** window, select **Storage Array > Premium Features > Enable**.

The **Select Feature Key File** dialog is displayed.

Note: The **Select Feature Key File** dialog filters for files with the `.key` extension.

2. Select the folder in which you placed the generated key file.
3. Select the appropriate key file, and then select **OK**.

The **Enable Premium Feature** dialog is displayed.

4. Select **Yes**.

The Remote Volume Mirroring premium feature is now enabled. The icon in the premium feature status area no longer displays a red slash. To further verify the status of the feature, select **Storage Array > Premium Features > List**.

5. Repeat step 1 through step 4 for the secondary storage array and any additional storage array that will participate in the Remote Volume Mirroring Feature.

6. Has the Remote Volume Mirroring feature been activated on both primary and secondary storage arrays?
 - **No** – Go to “Activating the Remote Volume Mirroring Feature” on page 137.
 - **Yes** – Go to “Creating Mirror Relationships” on page 142.

Activating the Remote Volume Mirroring Feature

Activating the Remote Volume Mirroring feature prepares the storage array to create and configure mirror relationships. After the feature is activated, the secondary ports for each controller are reserved and dedicated to Remote Volume Mirroring feature usage. In addition, a mirror repository volume is automatically created for each controller in the storage array.

If the current status of the Remote Volume Mirroring feature is Enabled/Deactivated, complete the following to activate the premium feature.

Important: The Free Capacity node or Unconfigured Capacity node that you select for the mirror snapshot repository volumes must have a total of 256-MB capacity available. Two mirror repository volumes are created on this capacity, one for each controller. You must upgrade any existing mirror repository volume with a capacity of 4 MB to 128 MB. For more information, refer to “Upgrading Mirror Repository Volumes” on page 139.

1. From the **Array Management** window, select **Storage Array > Remote Volume Mirroring > Activate**.

The **Activate Remote Volume Mirroring - Introduction** dialog is displayed.

Note: Select **Back** at any time to change information in previous dialogs.

2. Decide where the mirror repository volumes will reside. Select one of the following.
 - **Use free capacity on existing volume group** – If this option is selected, then a corresponding volume group must be selected as well.
 - a. Select a volume group.
 - b. Select **Next**.
 - c. Go to step 6.

- **Create a new volume group using unconfigured capacity** – If this option is selected, more information must be collected about the new volume group. Go to step 3.
3. In the **Create New Volume Group** dialog, select the RAID level for the new volume group.
 4. Under the drive selection choices, select one of the following.
 - **Automatic** – The drives are chosen automatically according to available capacity.
 - **Manual** – This option gives the user the option to specify which drives contain the mirror repository volume.

Note: If the RAID level for the new volume group is RAID 1, then the number of drives must be an even number. If an odd number of drives is selected, an error message will display when **Apply** is selected.

Note: RAID 0 is not an option, therefore, the number of drives selected must always be greater than one for all RAID levels.

Note: RAID 3 and 5 require a minimum of 3 drives.

Note: RAID 6 configurations require a minimum of 5 drives.

- a. Select the desired drives. Multiple drives can be selected by pressing **Ctrl+click**.
 - b. Select **Apply**.
5. Select **Next**.
 6. Review the information in the **Preview** dialog, then select **Finish** if all information is satisfactory.
 7. Is the storage array password protected?
 - **Yes** – Go to step 8.
 - **No** – Go to step 9.
 8. Enter the password and select **OK**.

9. Review the information in the Completed dialog, then select OK to finish the Remote Volume Mirroring feature activation.

The Remote Volume Mirroring premium feature is now activated. The icon in the premium feature status area changes from grey to green.

10. Repeat step 1 through step 9 for the secondary storage array and any additional storage array that will participate in the Remote Volume Mirroring Feature.
11. Has the Remote Volume Mirroring feature been enabled on both primary and secondary storage arrays?
 - **No** – Go to “Enabling the Remote Volume Mirroring Feature” on page 136.
 - **Yes** – Go to “Creating Mirror Relationships” on page 142.

Upgrading Mirror Repository Volumes

An upgrade to the mirror repository volumes is required only if a previous version of the storage management software existed with the Remote Mirroring feature installed. If this version of the storage management software (version 10.1x) is the first time that the Remote Mirroring feature is to be activated, then the procedures in this section do not apply.

Use the Upgrade Mirror Repository Volumes Wizard to create new, larger mirror repository volumes on your storage system, to move any data from the existing mirror repository volumes to the new ones, and then to delete the old mirror repository volumes.

In previous versions of the storage management software, a maximum of 32 volumes could participate in mirror relationships. The current version of storage management software lets 128 volumes participate in mirror relationships. If your controller enclosure or controller-drive enclosure permits 128 volumes to participate in mirror relationships, you must upgrade any existing mirror repository volumes on the storage system to have a minimum capacity of 256 MB. This capacity upgrade lets the previously mentioned controller enclosure or controller-drive enclosure have more than 32 volumes participating in mirror relationships.

Important: If there is not enough free capacity or unconfigured capacity available on your storage system, you cannot upgrade your mirror repository volumes. The node that you select must have a minimum of 256 MB of available capacity.

1. **Select Storage Array >> Remote Volume Mirroring >> Upgrade Mirror Repository Volumes.** The Upgrade Mirror Repository Volumes - Introduction dialog appears.

Important: The Free Capacity node or Unconfigured Capacity node that you select for the mirror repository volumes must have a total of 256-MB capacity available. Two mirror repository volumes are created for this capacity, one for each controller. You must upgrade any existing mirror repository volumes with a capacity of 4 MB to 128 MB.

2. Select one of the following three options based on where the upgraded mirror repository volumes will reside.
 - Free capacity on same drive group (create individual volume) – Go to step 3.
 - Free capacity on different drive group (create individual volume) – Go to step 3.
 - Unconfigured capacity (create new volume group) – Go to step 4.
3. Create an individual volume in the associated drive group.
 - a. Select a drive group (volume group).
 - b. Click **Next**.
 - c. Go to step 7.

Important: If the RAID level for the new drive group is RAID 1, then the number of drives must be an even number. If an odd number of drives is selected, an error message appears when you click Apply. RAID 0 is not an option; therefore, the number of drives selected must always be greater than one for all RAID levels. RAID 3 and RAID 5 configurations require a minimum of three drives. RAID 6 configurations require a minimum of five drives.

Note: If you need more information about the selections in Create New Volume Group dialog, click the Help button to launch the online help topic specific to the dialog.

4. In the Create New Volume Group dialog, select the RAID level for the new drive group.
5. Under the drive selection choices, select one of these options:

- **Automatic** – Choose this option to select the drives automatically according to available capacity. If chosen, click **Next**, and go to step 7.
 - **Manual** – Choose this option for more customization. Go to step 6.
6. You must specify which drives contain the mirror repository volume.
 - a. Select the desired drives. (To select multiple drives, press Ctrl-click.)
 - b. To move the drives to the Selected Drives area, click **Add**.
 - c. Click **Next**.
 7. Review the information in the Preview dialog, and click **Finish** if all information is satisfactory.

Note: The upgrade process might take a minute or longer to complete, depending on the current workload of the storage system. Any I/O activity to the storage system might be delayed during the upgrade process.

8. Are you asked for a password?
 - **Yes** – Enter the password, and click **OK**. Go to step 9.
 - **No** – Go to step 9.
9. Review the information in the Completed dialog, and click **OK**.

The mirror repository volumes are now upgraded, and the storage system can support up to 64 mirror relationships, depending on the model of the controller enclosure or the controller-drive enclosure that you are using.
10. Repeat step 1 through step 9 for the secondary storage system and any other storage system that contains mirror repository volumes that must be upgraded.

Creating Mirror Relationships

The Remote Volume Mirroring premium feature should be enabled and activated prior to creating mirror relationships. After the feature is enabled and activated, the secondary host ports on the storage arrays are reserved for data synchronization between a primary and secondary volume of a mirror relationship, and mirror relationships can be created.

This subsection will present information for the following topics.

- How to create volumes to be a primary or secondary volume in a mirror relationship, if not already created.
- How to use the **Create Remote Volume Mirror Wizard** to create mirror relationships.

Creating Volumes for Mirror Relationships

Before any mirror relationships can be created, volumes must exist at both the primary and secondary sites. The volume residing in the primary storage array is the primary volume. Similarly, the volume residing in the secondary storage array is the secondary volume.

If a primary volume does not exist, one will need to be created on the primary storage array.

If a secondary volume does not exist, one will need to be created on the secondary storage array.

Consider the following when creating the secondary volume:

- The secondary volume must be of equal or greater size than the primary volume.
- The RAID level of the secondary volume does not have to be the same as the primary volume.

For more information about creating volumes, see the **Array Management** window online help.

When adequate volumes exist at both sites, mirror relationships can be created using the **Remote Volume Mirror Wizard**.

Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard

The **Create Remote Volume Mirror Wizard** allows a user to create a mirror relationship between a primary volume and a secondary volume.

1. Before beginning the wizard, ensure the following requirements are met.
 - The Remote Volume Mirroring premium feature is enabled on both storage arrays in which the primary and secondary volumes will reside.
 - The Remote Volume Mirroring feature is activated on both arrays.
 - The storage arrays are connected through a proper Fibre Channel fabric configuration.
 - The secondary volume is RAID level 0, 1, 3, 5, or 6, and is of equal or greater size than the primary volume.
 - The RAID level of the secondary volume does not have to be same as the primary volume.
2. In the **Logical/Physical View** of the **Array Management** window, select the volume that will be the primary volume in the mirror relationship.
3. Select **Volume > Remote Volume Mirroring > Create**.

The **Create Remote Volume Mirror** dialog is displayed stating requirements similar to those mentioned at the beginning of this section.
4. Select **Next**.
5. Perform the procedure in the **Select Remote Storage Array** dialog.
 - a. Select the proper storage array.
 - b. Select **Next**.
6. Select the volume that will be the secondary volume in the mirror relationship.
7. Select **Next**.
8. Select the Synchronization Priority Level.

Note: For more information about the Synchronization Priority Level, see “Changing the Synchronization Settings and the Order Resynchronization Option” on page 147.

9. Select **Finish**.

10. If the remote storage array is password protected, proceed to step 10, otherwise, go to step 11.
11. Enter the password and select **OK**.
12. Create another mirror relationship?
 - **Yes** – Complete the following steps:
 - a. Select the volume that will become the next primary volume.
 - b. Select **Next**.
 - c. Repeat steps 4 through 10 for each additional mirror relationship.
 - **No** – Go to step 12.
13. Review the information in the **Completed** dialog, then select **OK** to complete the **Create Remote Volume Mirror Wizard**.

The primary and secondary volumes are now displayed with the mirror relationship icons, which will change slightly upon completion of the data synchronization. The amount of time for the data to be synchronized between the two volumes varies proportionally to the size of the volume being mirrored. After the synchronization is complete, the icons will change to the optimal state mirror relationship icons.

Note: To view all icons associated with the Remote Volume Mirroring feature and obtain a description of their purpose, refer to the **Array Management** window online help.

Figure 3-33 shows how the icons appear during and after data synchronization.

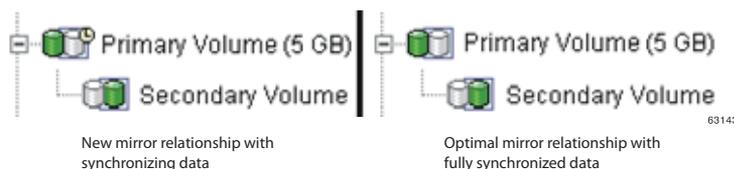


Figure 3-33 Synchronizing (Left) and Optimal (Right) Mirror Relationship

14. The next step is maintaining the mirror relationship. Go to “Maintaining Mirror Relationships” on page 145.

Maintaining Mirror Relationships

Once a mirror relationship is active, certain attributes of the relationship can be modified at any time. This section describes how to view information about the mirror relationship, change the synchronization priority level, remove a mirror relationship, and delete volumes participating in a mirror relationship.

Viewing Storage Array Profile

The mirror relationship that exists between the primary volume and the secondary volume can be examined by various methods, including: the **Storage Array Profile**, the **Mirroring Properties** dialog, and the **View Associated Components** dialog.

The Storage Array Profile

The **Storage Array Profile** is the most efficient way to view information about any or all components of the storage array. Details for all volumes, including the primary, secondary, and mirror repository volumes can all be easily viewed through the **Storage Array Profile** simultaneously. The **Storage Array Profile** also contains specific information for components associated with mirror relationships.

Complete the following procedure to display a **Storage Array Profile**.

1. From the **Array Management** window select **View > Storage Array Profile**.

Note: To save the information shown in the storage system profile as a `*.txt` file, click the **Save as** button. The **Save Profile** dialog lets you save the information in the current tab, save the information from multiple tabs, or save the entire storage system profile.

2. Select the **Volumes** tab.
 - a. To view information about the mirror relationships that exist on this storage array, select the **Mirrors** tab.
 - b. To view information about the mirror repository volume, select the **Repositories** tab.

Note: All information shown in the **Storage Array Profile** can be saved as a `.txt` file by selecting the **Save as** button. The **Save Profile** dialog that is displayed will have the option of saving the current window, multiple windows, or the entire **Storage Array Profile**.

Viewing the Mirroring Properties Dialog

The **Mirroring Properties** dialog displays all the physical characteristics of a single volume in the mirror relationship. The dialog displays the same information as the **Storage Array Profile** for the selected volume, but is specific to that volume. In addition, the synchronization progress is displayed if the selected volume is synchronizing data with another volume in a mirror relationship.

Note: Use the following procedure to view the synchronization progress of a recently created mirror relationship.

Complete the following procedure for each volume that needs to be examined in the mirror relationship.

1. Select the primary or secondary volume.
2. Select **Volume > Properties**.
3. Select the **Mirroring** tab.

The View Associated Components Dialog

The **View Associated Components** dialog provides a more graphical representation of the volumes participating in the mirror relationship. In addition, more detail is given for all components, rather than only the selected volume.

Complete the following procedure to view all associated components in a mirror relationship, including primary, secondary, and mirror repository volumes.

1. Select the primary or secondary volume in a mirror relationship.
2. Select **View > Associated Components**.

The **View Associated Components** dialog is displayed.

Changing the Synchronization Settings and the Order Resynchronization Option

The Synchronization Priority Level of a mirror relationship defines the amount of system resources used to synchronize the data between the primary and secondary volumes of a mirror relationship. If the highest priority level is selected for a mirror relationship, the data synchronization uses a high amount of system resources to perform the full synchronization, but may decrease performance for host I/O, including other mirror relationships. Conversely, if the lowest synchronization level is selected, there is less impact on complete system performance, but the full synchronization may be slower.

To change the Synchronization Priority Level for a mirror relationship, complete the following procedure.

1. In the **Logical/Physical View** of the **Array Management** window, select a primary volume of a mirror relationship.
2. Select **Volume > Remote Volume Mirroring > Change > Synchronization Priority**.
The **Change Synchronization Priority** dialog is displayed. The primary volume chosen in step 1 is selected by default in the **Select Volumes** selection area.
3. Select one or more volumes using the **Select All** button, **Ctrl+click**, or **Shift+click**.
4. Select the synchronization priority level by moving the slider bar. The five levels are Lowest, Low, Medium, High, and Highest. All selected volumes will change to the same synchronization priority level.
5. Select **OK**.
6. Select **Yes** when the **Confirmation** dialog is displayed.
7. Select **OK** when the **Completed** dialog is displayed.

Suspending a Mirror Relationship

Use the Suspend option to stop data transfer between a primary volume and a secondary volume that are participating in a mirror relationship without removing the mirror relationship.

Suspending a mirror relationship lets you control when the data on the primary volume and the secondary volume should be synchronized. This action helps to reduce any performance impact to the host application that might occur while any changed data on the primary volume is copied to the secondary volume.

When a mirror relationship is in a Suspended state, no attempt is made to contact the secondary volume. Any writes to the primary volume are persistently logged in the mirror repository volumes. After the mirror relationship is resumed, only the modified regions of the primary volume are written to the secondary volume. You can use the Suspend option to back up a secondary volume.

Important: Suspending a mirror relationship removes any Needs Attention status resulting from communication failure between the primary volume and the secondary volume.

Important: Possible loss of data access – If the selected mirrored pair is part of a write consistency group, you automatically suspend all of the mirrored pairs in the write consistency group. Use the command line interface to resume single write-consistent mirrored pairs. For more information about using the command line interface, refer to the online help topics in the Enterprise Management Window. For more information about write consistency groups, see “Write Consistency Mode” on page 85.

Important: Any data that is written to the primary volume is logged while the mirror relationship is suspended. This data is automatically written to the secondary volume when the mirror relationship is resumed. A full synchronization is not required.

The state of the remote mirror remains suspended until you use the **Resume** option to resume synchronization activity.

1. In the Logical/Physical View of the Array Management Window, select a primary volume of a mirror relationship.
2. Select **Volume >> Remote Volume Mirroring >> Suspend**. The Suspend Mirrored Pair dialog appears.

3. Select one or more mirror relationships to suspend. To make multiple selections, click **Select All**, press **Ctrl-click**, or press **Shift click**.
4. Click **Suspend**.

Important: Possible loss of data access – If the selected mirrored pair is part of a write consistency group, you automatically suspend all of the mirrored pairs in the write consistency group. For more information about using the command line interface, refer to the online help topics in the Enterprise Management Window. For more information about write consistency groups, see “Write Consistency Mode” on page 85.

5. Type **Yes**, and click **OK**.
The **Suspend Mirrored Pair - Progress** dialog appears while the selected mirrored volumes are suspended. After all of the selected mirrored pairs have been suspended, the **OK** button becomes active.
6. Click **OK**.
The **Array Management Window** appears, and the primary volume and the secondary volume are shown with an **Optimal/Suspended** status.

Suspending Mirrors When Operating in Asynchronous Mode

You can put an asynchronous group of mirrors into a consistent state at the remote site. When the volumes are consistent, you can make snapshots of the volumes. This process makes sure that database applications are always in a recoverable state in the event of a site failure.

If you have an asynchronous group of mirrors, the group must first become synchronized between the primary site and the secondary site before a valid snapshot can be made. This statement is especially true when multiple volumes are involved.

This procedure explains how to perform these actions:

- Force an asynchronous mirror pair into Synchronous mode
- Make sure that the mirrors are synchronized with each other
- Suspend the mirrors
- Create a snapshot

- Resume the mirrors
- Reset the mode of the mirror pair to Asynchronous mode

Use this procedure to suspend mirrors in Asynchronous mode:

Note: These commands are executed using the ISSM-EE command line interface (smeecli). For more information on the smeecli, refer to the *ISSM-EE Command Line Interface and Script Commands Programming Guide*. Table 3-3 lists the variable parameter definitions used in these commands.

1. Type the following command, and press Enter.

```
set remoteMirror (localVolume  
[user-label] | localVolumes [user-label-list])  
writeMode=synchronous;
```

2. Type the following command, and press Enter.

```
check remoteMirror localVolume  
[user-label] optimalStatus timeout=time-interval;
```

3. Type the following command, and press Enter.

```
suspend remoteMirror (primary [user-label] |  
primaries [user-label-list]) [writeConsistency=boolean];
```

4. Type the following command, and press Enter.

```
recreate snapshot (volume [user-label] |  
volumes [user-label-list]);
```

5. Type the following command, and press Enter.

```
resume remoteMirror (primary [user-label] |  
primaries [user-label-list]) [writeConsistency=boolean];
```

6. Type the following command, and press Enter.

```
set remoteMirror (localVolume  
[user-label] | localVolumes [user-label-list])  
writeMode=asynchronous;
```

Table 3-3 Variable Definitions

Parameter	Definition
<i>user-label</i>	The name of the volume on which the operation is being made. This variable is usually enclosed in double quotation marks around the parameter in square brackets. If you use this variable in a script, the double quotation marks might need to be escaped. For example, ["Database Log Volume1"] or ["Database Log Volume1\"].
<i>user-label-list</i>	A list of volumes, separated by spaces, using the same convention as <i>user-label</i> . For example, ["Database Log Volume1" "Database Log Volume2"].
<i>time-interval</i>	The value (in minutes) that the <code>check remoteMirror</code> command waits before timing out.

An example of the `smeecli` (command line) script that uses these constructs follows:

```
set remoteMirror localVolume ["database log volume"]
writeMode=synchronous;
check remoteMirror localVolume ["database log volume"]
optimalStatus timeout=1;
suspend remoteMirror localVolume ["database log volume"];
recreate snapshot volume ["database log volume"];
resume remoteMirror primary ["database log volume"];
set remoteMirror localVolume ["database log volume"]
writeMode=asynchronous;
```

Typically, error checking is performed, which prevents a single `smeecli` script from being run in its present condition. Instead, use a command script to invoke the `smeecli` commands individually, and verify their status before proceeding to the next command.

An example of such a script used to suspend a primary site follows:

```
#!/bin/sh
#
# Syntax: rmsuspendasync.sh <rmVolume> This command is used to
# suspend a single rm volume
# This requires the primary volume to be referenced on the
# primary storage system
#
# Return codes:
# 0 - Success
# 1 - Suspend was not successful
# 2 - Mirror was already suspended
# 9 - Script error
#
# This is a sample script and should be used as a guideline
# only.
# Additional parameters such as multiple volumes or consistency
# will require modifications to the script and potentially the
# logic.
if [ "${1}" == "" ]; then
echo "Syntax: rvmsuspend.sh <rmVolume>"
echo "Where rmVolume is the primary volume"
else
SMEECLI=~ /opt/smee/client/smeecli -n "D4000_Development" -c \
"suspend remoteMirror primary [\"${1}\"];\"`
echo ${SMEECLI} | grep -q "^. * completed successfully."
if [ $? == 0 ]; then
echo "RM volume ${1} suspended successfully"
exit 0
fi
echo ${SMEECLI} | grep -q "attempted to suspend an already
suspended mirror."
if [ $? == 0 ]; then
echo "RM volume ${1} already suspended"
exit 2
fi
echo ${SMEECLI} | grep -q "Script execution halted due to error."
if [ $? == 0 ]; then
echo "RM volume ${1} suspend unsuccessful"
exit 1
fi
echo "We should never have gotten to this point. Script error"
exit 9
fi
```

Resuming a Mirror Relationship

Use the Resume option to restart data transfer between a primary volume and a secondary volume that are participating in a mirror relationship, after the mirror has been suspended or unsynchronized.

This option lets you control when the data on the primary volume and the secondary volume should be resynchronized. This option helps reduce any performance impact to the host application that might occur while any changed data on the primary volume is copied to the secondary volume.

After the mirrored pair is resumed, only the regions of the primary volume that have changed since the mirrored pair was suspended are written to the secondary volume.

Important: Possible loss of data access – If the selected mirrored pair is part of a write consistency group, you automatically resume all of the mirrored pairs in the write consistency group. For more information about using the command line interface, refer to the online help topics in the Enterprise Management Window. For more information about write consistency groups, see “Write Consistency Mode” on page 85.

1. In the **Logical/Physical View** of the **Array Management Window**, select a primary volume of a mirror relationship.
2. Select **Volume >> Remote Volume Mirroring >> Resume**. The **Resume Mirrored Pair** dialog appears.
3. Click **Resume**. The **Resume Mirrored Pair - Confirmation** dialog appears.
4. To resume data transfer to the mirrored pairs that you selected, click **Yes**.

The Resume Mirrored Pair - Progress dialog appears while data transfer is resumed to the selected mirrored pairs. After data transfer has been resumed to the mirrored pairs, the **OK** button becomes available.

5. Click **OK**.
6. The **Array Management Window** appears, and the primary volume and the secondary volume are shown with an **Optimal** status or a **Synchronizing** status.

Removing Mirror Relationships

Removing a mirror relationship between a primary and secondary volume does not affect any of the existing data on either volume. The link between the volumes is removed, but the primary volume still continues normal I/O operation. This action can be performed for backup routines, particularly from the secondary volume. A mirror relationship between the two volumes can be re-created unless one of the volumes is deleted.

To remove a mirror relationship between two volumes, complete the following procedure.

1. From the **Array Management** window, select a local primary volume or local secondary volume of a mirror relationship.
2. Select **Volume > Remote Volume Mirroring > Remove Mirror Relationship**.
The **Remove Mirror Relationship** dialog appears displaying all mirror relationships associated with this storage array.
3. Select one or more mirror relationships to remove using the **Select All** button, **Ctrl+click**, or **Shift+click**.
4. Carefully review the information presented in the **Confirmation** dialog.
5. Select **Yes**.

Changing the Write Mode

When the Synchronous write mode is selected, the controller on the primary storage system waits until data has been successfully copied to the secondary storage system before sending an I/O completion indication back to the host system. If the Asynchronous write mode is selected, the controller on the primary storage system sends an I/O completion indication to the host system before the data has successfully been copied to the secondary storage system.

1. From the **Array Management Window**, select a local primary volume of a mirror relationship.
2. Select **Volume >> Remote Volume Mirroring >> Change >> Write Mode**.

The Change Write Mode dialog appears.

3. Select one or more mirrored pairs.

To make multiple selections, click **Select All**, press **Ctrl-click**, or press **Shift-click**.

4. Select one of the following modes:

- **Synchronous** – Synchronous mode offers the best chance of full data recovery from the secondary storage system in the event of a disaster; however, this data recovery happens at the expense of host I/O performance. The Synchronous mode is the preferred mode of operation. You have completed this procedure.
- **Asynchronous** – Asynchronous mode offers faster host I/O performance. This mode, however, does not guarantee that data was successfully written to the secondary volume before indicating a successful write to the host system. You have completed this procedure.
- **Write Consistency** – Go to step 5.

Note: Selecting this option might impact host I/O performance.

5. To activate Write Consistency mode:

- a. Select **Asynchronous**.
- b. Select the **Add to write consistency group** check box. This action makes sure that writes to the secondary storage system are completed in the same order as on the primary storage system.

For more information, see “Write Consistency Mode” on page 85.

Deleting Primary or Secondary Volumes

Deleting a volume participating in a mirror relationship removes the mirror relationship and completely deletes the volume from the storage array. The mirror relationship cannot be redefined until a new volume is created or an alternate volume is chosen to replace the deleted volume.

To delete a primary or secondary volume from a mirror relationship, complete the following procedure.

Note: This action permanently removes data from the selected volume.

1. From the **Array Management** window, select the primary or secondary volume that will be deleted.
2. Select **Volume > Delete**.

The following information is displayed in the **Delete Remote Volume Mirror** dialog.

Caution: Carefully review the following information before proceeding with step 3.

Deleting this volume will destroy all data on it, remove the mirror relationship that exists between it and its associated volume in the mirror pair, and increase the Free Capacity on the volume group.

NOTE: The Data on the associated volume will not be deleted.

Stop all I/O and unmount any file systems on this volume before proceeding

IMPORTANT: If you see system error messages related to this deleted volume, you may need to reboot your host system to permanently remove any information about this volume.

Are you sure you want to continue?

3. Type **Yes** in the text area to delete the volume.
4. Select **OK**.

Deactivating and Disabling the Remote Volume Mirroring Feature

If no mirror relationships exist and the Remote Volume Mirroring feature is no longer required, then deactivating the feature will reestablish normal use of dedicated ports on both storage arrays and delete both mirror repository volumes.

The Remote Volume Mirroring feature can also be disabled. When the feature is in the Disabled/Active state, previously existing mirrors can still be maintained and managed, however, new mirror relationships cannot be created. When in the Disabled/Deactivated state, no Remote Volume Mirroring feature activity can occur.

Deactivating the Remote Volume Mirroring Feature

To change the Remote Volume Mirroring feature status from Enabled/Activated to Enabled/Deactivated or from Disabled/Activated to Disabled/Deactivated, complete the following procedure.

Caution: This procedure will delete all mirror repository volumes on the current storage array.

1. Verify that all mirror relationships have been removed. For more information, see “Removing Mirror Relationships” on page 154.
2. Select **Storage Array > Remote Volume Mirroring > Deactivate**.
3. Carefully review the information presented in the **Confirmation** dialog.
4. Select **Yes**.

Disabling the Remote Volume Mirroring Feature

To change the Remote Volume Mirroring feature status from Enabled/Deactivated to Disabled/Deactivated or from Enabled/Activated to Disabled/Activated, complete the following procedure.

1. From the **Array Management** window, select **Storage Array > Premium Features > Disable**.
2. Select **Remote Volume Mirroring**.
3. Select **OK**.
4. Carefully review the information presented in the **Confirmation** dialog.
5. Select **Yes**.

Disaster Recovery

This section describes the actions to take to prepare for a disaster and the actions to take if a hardware component or complete site failure should occur.

As businesses require the ability to access critical data round-the-clock with little or no downtime, system administrators are required to ensure that critical data is safeguarded against potential disasters whether natural or man-made.

In preparing for a potential disaster, system administrators need to develop a disaster recovery plan that describes the procedures devised to prevent and prepare for disasters, and proposed actions for responding to and recovering from disasters if they occur.

- “Planning for Disaster Recovery” on page 158
- “Recovering from Switch Failures” on page 160
- “Recovering from Storage System Failures” on page 161
- “Reversing the Roles of Primary and Secondary Volumes” on page 164

Planning for Disaster Recovery

This section provides information on the steps that should be undertaken in preparation for a possible disaster or catastrophic failure. Preparation steps discussed include backing up critical data prior to and during Remote Volume Mirroring, using host clustering software for failover protection, and preparing the secondary storage array that will participate in the Remote Volume Mirror.

The Remote Volume Mirroring checklist should be completed to ensure that your primary and secondary storage arrays are prepared correctly.

Backing Up Critical Data

Backing up critical data regularly is vital insurance against disasters or catastrophic failure. Backups should be performed regardless of whether the Remote Volume Mirroring feature is in use.

A backup of all critical data on both the primary and secondary storage arrays participating in Remote Volume Mirroring should be performed prior to creating a

Remote Volume Mirror. Periodic backups of the primary storage array should be carried out after Remote Volume Mirrors are in use.

Backups to tape or disk can be made while I/O continues using the Snapshot Volume premium feature, which captures a point-in-time image of a volume. For more information on using snapshot volumes, see “Snapshot Volume” on page 95 or the **Array Management** window online help.

Preparing the Secondary Storage Array

When preparing for a potential disaster, ensure the secondary storage array, to which the data is being mirrored, is structured in the same manner as the primary storage array.

Prior to creating a Remote Volume Mirror, you will need to prepare your secondary storage array to ensure that storage partitions are defined and ready to function if secondary storage array volumes are promoted to primary roles.

Defining Storage Partitioning

The Storage Partitioning premium feature allows you to define which hosts have access to which volumes on the storage array and at what logical unit number. When configuring your secondary storage array, replicate the storage partitions that were previously defined on the primary storage array. This will ensure that after a role reversal is carried out, hosts will be able to access their required data.

For more information on Storage Partitioning, see “Using Other Premium Features and Remote Volume Mirroring” on page 94 or the **Array Management** window online help.

Ensuring Failover Protection Using Host Clustering Software

Host failover clustering is a process in which the operating system and application software work together to provide continuous availability in the event of an application failure, hardware failure, or operating-system error. Software products that support remote failover may refer to global or geographic cluster management.

The Remote Volume Mirroring feature can be used in conjunction with remote host failover to provide complete redundancy of host systems and storage arrays. If a disaster occurs at the primary storage array, hosts at the secondary storage array will take over processing and initiate a role change to the secondary storage array.

Remote failover is not a requirement for a disaster recovery plan. Depending on the requirements for how quickly operations must be functioning at Optimal status again, mirroring the storage arrays and depending on manual startup of the remote hosts may be acceptable.

Disaster Preparation Checklist

The checklist in Table 3-4 should be completed in preparation for disaster recovery. Complete the steps listed to ensure access to mirrored data is possible in the event of a disaster or catastrophic failure.

Table 3-4 Disaster Preparation Checklist

Task Description	Validation Activity
1. Ensure that all critical data has been backed up on both the primary and secondary storage arrays.	<input type="checkbox"/> Completed.
2. Ensure that Storage Partitioning is in place on the secondary storage array. This will reduce any lag time when promoting a secondary volume to the primary volume role.	<input type="checkbox"/> Completed.
For more information on setting up Storage Partitioning, see the Array Management window online help.	
3. If possible, ensure that the your host clustering software is configured to allow host failover between the primary storage array and secondary storage array participating in Remote Volume Mirroring.	<input type="checkbox"/> Completed.
For more information, see your host clustering software documentation.	

Recovering from Switch Failures

The Campus Configuration is the only configuration where each switch is a single point of failure. In all other configurations, if a switch fails, at least one other switch is configured to be an alternate path for all I/O. The single switch at each site is not only responsible for the communication between the two sites for the Remote Volume Mirroring feature, but for all host to storage array I/O as well.

If a switch fails in the Campus Configuration, the system administrator should decide whether or not to do a complete site failover. This decision should be based on the immediate availability of a backup switch or the time it will take to receive a replacement.

If a replacement switch is not readily available, one of the following temporary scenarios will need to occur until the switch can be replaced.

- Direct connection between hosts and storage arrays – This will bypass the switch and allow normal I/O to continue. All mirror relationships will be suspended until a normal configuration is resumed. Since the storage arrays have a finite number of host connections, all hosts may not be able to access the storage array.
- Entire Site Failover – The switch at the primary site has failed, however, the switch at the secondary site is still functional, which allows all hosts to access the switch and storage arrays. Mirror relationships will be suspended until the primary storage array is recovered. If a complete site failover is needed to continue normal operations, see “Entire Site Failover Procedure” on page 162.

Recovering from Storage System Failures

This section provides procedures for recovering storage systems where one of the following situations has occurred:

- The primary storage array is damaged or destroyed.
- Entire primary site is damaged or destroyed.

Both situations require a complete site failover from the primary site to the secondary site so that business operations can continue. For a secondary site failure, no site failover is necessary. If the primary volume is not damaged, then a full synchronization will occur when the site is recovered. If the primary storage array is damaged, follow the procedure in “Rebuilding a Damaged Site” on page 163.

Note: The procedures in this section should be completed in the order presented to ensure that the affected site or storage array is recovered successfully.

Entire Site Failover Procedure

Entire site failover is necessary when an event occurs that inhibits normal operation to the primary site for an extended period of time.

1. Manually reverse the role of the secondary volume of the mirror relationship to the primary role. For more information, see “Reversing the Roles of Primary and Secondary Volumes” on page 164.
2. Have all hosts at the secondary site been properly configured for host failover?
 - **Yes** – Go to step 4.
 - **No** – Go to step 3.
3. Reinstall host cluster software and reconfigure settings as needed. For more information about installing and configuring host cluster software, refer to the documentation provided with the software.
4. Has Storage Partitioning been pre-configured so that proper hosts can access appropriate volumes?
 - **Yes** – Go to step 6.
 - **No** – Go to step 5.
5. Configure the storage partitioning so that host to volume access is identical to the primary site configuration. For more information about setting up storage partitioning, see “Using Other Premium Features and Remote Volume Mirroring” on page 94.
6. Resume normal operation with the secondary site acting as a new, fully-functioning primary site.

The entire site failover is complete. The primary site must now be reconfigured and the Remote Volume Mirroring environment must be re-created. The next step will depend on the resulting status of the storage array from the disaster.
7. Is the storage array and all the contained data recovered and fully operational?
 - **Yes** – Go to “Re-creating a Mirror Relationship” on page 163.
 - **No** – Go to “Rebuilding a Damaged Site” on page 163.

Rebuilding a Damaged Site

If a disaster occurs to the primary or secondary site that renders the storage array or all data on the storage array irreparable, then the site needs to be rebuilt with new equipment.

Complete the following procedure to rebuild a damaged site:

1. Obtain new storage arrays and any hosts and equipment needed to properly use the Remote Volume Mirroring feature. It is recommended that the new equipment have equal or greater hardware and software specifications than the damaged equipment.
2. Depending on the configuration, complete the switch zoning and hardware installation procedure as presented in “Hardware and Software Installation” on page 108 for this site only.
3. Create a volume to replace the mirrored volume damaged on the original storage array. The new volume should be identical to the specifications of the previously damaged volume.
4. Ensure storage partitioning is properly defined so that it may take over normal operation from the secondary site. For more information on setting up storage partitioning, see “Using Other Premium Features and Remote Volume Mirroring” on page 94.
5. Ensure the host cluster software is properly configured so that the hosts can take over I/O from the secondary site hosts. For more information about installing and configuring host cluster software, refer to the documentation provided with the software.
6. The damaged site is now ready to support mirror relationships. Go to “Re-creating a Mirror Relationship” on page 163.

Re-creating a Mirror Relationship

After the damaged site is back online and properly configured, mirror relationships can be resumed by completing the following procedure.

1. From the undamaged site, remove all mirror relationships that existed prior to the disaster.
2. From the undamaged site, create a mirror relationship using the volume on the recovered primary site as the secondary volume. For more information, see “Maintaining Mirror Relationships” on page 145.

3. Ensure storage partitioning is properly defined on the recovered primary site so that it may take over normal operation from the secondary site. For more information on setting up storage partitioning, see “Using Other Premium Features and Remote Volume Mirroring” on page 94.
4. Ensure the host cluster software is properly configured so that the hosts at the recovered primary site can take over I/O from the secondary site hosts. For more information about installing and configuring host cluster software, see the documentation provided with the software.
5. Perform a role reversal so that the recovered primary site now possesses the active primary volume, and the secondary volume now exists on the secondary site. For more information, see “Reversing the Roles of Primary and Secondary Volumes” on page 164.

The Remote Volume Mirroring configuration is now ready.

Reversing the Roles of Primary and Secondary Volumes

A role reversal is the act of promoting the secondary volume to be the primary volume within the mirrored volume pair, and demoting the primary volume to be the secondary volume.

A role reversal may be performed using one of the following methods:

- Changing a secondary mirrored volume to a primary volume – This option promotes the selected secondary volume to become the primary volume of the mirrored pair and would be used when a catastrophic failure has occurred. For step-by-step instructions, see “Changing a Secondary Volume to a Primary Volume” on page 165.
- Changing a primary mirrored volume to a secondary volume – This option demotes the selected primary volume to become the secondary volume of the mirrored pair and would be used during normal operating conditions. For step-by-step instructions, see “Changing a Primary Volume to a Secondary Volume” on page 165.

Role reversals may also be performed using the `set` command in the Script Editor or the command line interface. For more information, see the **Enterprise Management** window online help.

Changing a Secondary Volume to a Primary Volume

A secondary volume is typically promoted to a primary volume role when a catastrophic failure has occurred to the storage array that contains the primary volume, and the secondary volume needs to be promoted so that hosts can access data and business operations can continue.

This option is not available unless the Remote Volume Mirroring premium feature has been activated.

Note: When the secondary volume becomes a primary volume, any hosts that are mapped to the volume through a volume-to-LUN mapping will now be able to read or write to the volume.

If a communication problem between the secondary and primary sites prevents the demotion of the remote primary volume, an error message is displayed. However, you are given the opportunity to proceed with the promotion of the secondary volume, even though this will lead to a dual-primary condition.

To promote a secondary volume to a primary role:

1. Select the secondary volume in the **Logical** view, then select **Volume > Remote Mirroring > Change > Role to Primary**.
The **Change to Primary** dialog is displayed.
2. Select **Yes**. The secondary volume is promoted to be the primary volume role in the Remote Volume Mirror.

After the controller owner of the primary volume can be contacted, the primary volume is automatically demoted back to be the secondary role in the mirror relationship.

Changing a Primary Volume to a Secondary Volume

Changing a primary volume to a secondary role is used for role reversals during normal operating conditions. You can also use this option during a **Recovery Guru** procedure when a dual-primary condition occurs.

This option is not available unless the Remote Volume Mirroring premium feature has been activated.

Note: Any hosts that are accessing the primary volume through a volume-to-LUN mapping will no longer be able to read or write to the volume. When the primary volume becomes a secondary volume, only remote writes initiated by the primary controller will be written to the volume.

If a communication problem between the primary and secondary sites prevents the promotion of the secondary volume, an error message is displayed. However, you are given the opportunity to proceed with the demotion of the primary volume, even though this will lead to a dual-secondary condition.

If the selected primary volume has associated snapshot volumes, demoting this volume to a secondary role will cause the associated snapshot volumes to fail.

To demote a primary volume to a secondary role:

1. Select the primary volume in the **Logical** view, then select **Volume > Remote Mirroring > Change > Role to Secondary**.

The **Change to Secondary** dialog is displayed.

2. Select **Yes**. The primary volume is demoted to be the secondary volume in the Remote Volume Mirror.

Once the controller owner of the secondary volume can be contacted, the secondary volume is automatically promoted back to the primary volume in the mirror relationship.

Troubleshooting

This chapter provides information about troubleshooting some of the more commonly encountered problems when using the Remote Volume Mirroring feature.

Locate the section that directly relates to your problem or consult the section describing general troubleshooting techniques. If your problem is not discussed in any of the sections in this chapter, contact technical support.

Testing Communication

Use the Test Communication dialog to diagnose possible communication problems between a primary volume and a secondary volume that are participating in a mirror relationship.

After the communication test has been completed, the dialog shows either a Passed status, a Warning status, or a Failed status, together with values for these items:

- Test data round-trip time
- Maximum allowed round-trip time
- Average round-trip time

Important: If the communication test returns a Failed status, the test continues to run after you close the Test Communication dialog until communication between the mirrored pair is restored. If you have enabled automatic resynchronization, the mirrored pair automatically synchronizes. Otherwise, select **Volume >> Remote Volume Mirroring >> Resume** to manually start resynchronization.

1. In the **Logical/Physical View**, select a primary volume or a secondary volume.
2. Select **Volume >> Remote Volume Mirroring >> Test Communication**.

The Communication Test dialog appears. If the communication test returns a Passed with Caution status, the primary volume and the secondary volume are communicating correctly; however, the test data round-trip time is approaching the maximum time allowed. This status indicates that you need to take some type of action to fix this problem. If the communication test returns a Failed status, the reason for the failure is indicated, and you are directed to refer to **Recovery Guru** to correct the problem.

3. Review the information shown in the Communication Test dialog, and click **OK** to return to the **Array Management Window**.

General Troubleshooting

This subsection describes some of the typical problems that may be encountered when using the Remote Volume Mirroring premium feature. The information is presented in order from the time the feature is enabled and activated, to problems that may be encountered completing typical tasks when using the Remote Volume Mirroring feature.

Information about the probable cause and recommended resolution for each problem is provided. If your problem is not addressed in this section, see the remaining subsections in this section before contacting technical support.

- **The Remote Volume Mirroring menu options in the Array Management window are greyed out and unavailable.**

Cause:

The Remote Volume Mirroring premium feature is not enabled or activated.

Resolution:

1. Enable the Remote Volume Mirroring premium feature. A feature key file is required to enable premium features. If one has not been supplied, contact your storage supplier.
2. Activate the Remote Volume Mirroring feature.

For step-by-step instructions on enabling and activating the Remote Volume Mirroring feature, see “Enabling and Activating the Remote Volume Mirroring Feature” on page 134 or the **Array Management** window online help.

- **The Remote Volume Mirroring icon in the premium feature status area is greyed out with a red line.**

Cause:

The Remote Volume Mirroring premium feature is not enabled or activated.

Resolution:

1. Enable the Remote Volume Mirroring premium feature. A feature key file is required to enable premium features. If one has not been supplied, contact your storage supplier.
2. Activate the Remote Volume Mirroring feature.

For step-by-step instructions on enabling and activating the Remote Volume Mirroring feature, see “Enabling and Activating the Remote Volume Mirroring Feature” on page 134 or the **Array Management** window online help.

- **When creating a Remote Volume Mirror, an error message listing the message "Error 158 – The operation cannot complete because the selected volume is not a valid candidate for remote volume mirroring" is received.**

Cause:

The secondary volume candidate was participating in a mirror relationship with a different volume on this storage array. The mirror relationship was successfully removed from the other storage array. However, a communication error prevented the successful removal of the mirror relationship on this storage array.

A volume on this storage array still has a mirror relationship with the secondary volume candidate.

Resolution:

1. Select the **View > Storage Array Profile** option, then select the **Volumes** tab and the **Mirrors** tab.
 2. Locate the volume that is still participating in the mirror relationship with the secondary volume candidate and remove the mirror relationship. For step-by-step instructions, see “Removing Mirror Relationships” on page 154 or the **Array Management** window online help.
 3. Re-create the intended Remote Volume Mirror using the original candidate volumes. For step-by-step instructions, see “Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard” on page 143 or the **Array Management** window online help.
- **The primary volume is showing a Synchronization-in-Progress status.**

Cause:

A full synchronization, between the primary and secondary volumes participating in a mirror relationship, is in progress. All data on the primary volume is being copied in its entirety to its associated secondary volume.

Resolution:

Wait for the full synchronization to complete. After it has finished copying data from the primary volume to the secondary volume, the mirrored volume pair transitions to a Synchronized status.

For more information on monitoring the status of a Remote Volume Mirror, see the **Array Management** window online help.

- **Both the primary and secondary volumes are displayed as being the primary volume.**

Cause:

A disaster or catastrophic failure has occurred on the storage array containing the primary volume, and the secondary volume has been promoted to the primary volume role.

The primary storage array is now operational, but unreachable because of a link failure. A forced promotion of the secondary volume has resulted in both the primary and secondary volumes viewing themselves in the primary volume role.

Resolution:

1. From the **Array Management** window of the primary site, select the primary volume.
2. Select **Volume > Remote Volume Mirroring > Remove Mirror Relationship**.

Note: The mirror relationship for the secondary volume will not be removed if there is a connection problem between the storage arrays. If the mirror relationship still exists at the secondary site, repeat steps 1 and 2 for the secondary volume from the secondary site **Array Management** window.

3. Re-create the mirror relationship by completing “Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard” on page 143.
- **Both the primary and secondary volumes are displayed as being the secondary volume.**

Cause:

In response to a dual primary volume condition, a role reversal was carried out and the primary volume was demoted to the secondary volume role.

The secondary storage array is operational, but unreachable because of a link failure. The forced demotion of the primary volume has resulted in both the primary and secondary volumes viewing themselves in the secondary volume role.

Resolution:

1. From the **Array Management** window of the primary site, select the local secondary volume.
2. Select **Volume > Remote Volume Mirroring > Remove Mirror Relationship**.

Note: The mirror relationship for the secondary site will not be removed if there is a connection problem between the storage arrays. If the mirror relationship still exists at the secondary site, repeat steps 1 and two for the secondary volume from the secondary site **Array Management** window.

3. Re-create the mirror relationship by completing “Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard” on page 143.
- **The primary volume/secondary volume is displayed as being a Missing Volume.**

Cause:

A physical hard drive on which a Remote Volume Mirror depends (primary volume, secondary volume) has been removed, and the volume is now displayed under the Missing Volumes node in the **Logical** view of the **Array Management** window.

Resolution:

- If missing volumes are detected because drives have been accidentally removed, or are detected as missing because of a loss of power to the drive enclosures, recovery of these volumes is possible by:
 1. Re-inserting the drives back into the drive enclosure.
 2. Ensuring that the drive enclosure's power supplies are properly connected to an operating power source and have an Optimal status.

For more information, see **Recovery Guru** or contact technical support.

- If the missing volume is no longer required, it may be deleted.

Note: Deleting a missing volume is a permanent action. Any associated volumes or volume-to-LUN mappings will also be deleted. For more information, see the **Recovery Guru** or the **Array Management** window online help.

- **The mirror relationship was removed successfully from the primary side, but not from the secondary side.**

Cause:

When removing a mirror relationship, a link failure occurred between the primary and secondary storage arrays, and an error message was displayed.

The **Array Management** window for the primary storage array indicates that the mirror relationship has been removed successfully, however, the **Array Management** window for the secondary storage array indicates that the mirror relationship still exists.

Resolution:

Open the **Array Management** window for the secondary storage array and remove the mirror relationship.

For step-by-step instructions, see “Removing Mirror Relationships” on page 154 or the **Array Management** window online help.

- **The mirror relationship was removed successfully from the secondary side, but not from the primary side.**

Cause:

When removing a mirror relationship, a link failure occurred between the primary and secondary storage arrays.

The **Array Management** window for the secondary storage array indicates that the mirror relationship has been removed successfully, however, the **Array Management** window for the primary storage array indicates that the mirror relationship still exists, with the primary volume showing an Unresponsive status.

Resolution:

Open the **Array Management** window for the primary storage array and remove the mirror relationship.

For step-by-step instructions, see “Removing Mirror Relationships” on page 154 or the **Array Management** window online help.

- **The primary volume is showing an Unsynchronized status.**

Cause:

The Remote Volume Mirror is in an Unsynchronized status because the data mirrored between the primary volume and secondary volume is no longer identical. Typical causes for this status include link errors, a failed primary volume or secondary volume, or a dual primary or dual secondary error condition.

While the Remote Volume Mirror is in an Unsynchronized status, no mirroring activity will take place.

Resolution:

Once the controller owner of the primary volume can communicate with the controller owner of the secondary volume and the volume is online and optimal, a full synchronization will automatically take place. When a full synchronization begins, the Remote Volume Mirror transitions to a Synchronization-in-Progress status.

For information on monitoring the status of a Remote Volume Mirror, see the **Array Management** window online help.

- **Error message indicates that the Remote Volume Mirroring premium feature is out of compliance.**

Cause:

The Remote Volume Mirroring premium feature is out of compliance. This normally occurs if a volume group with volumes participating in Remote Volume Mirroring, has been moved into the primary storage array and one of the following conditions now exists:

- The Remote Volume Mirroring premium feature is not enabled on the selected storage array.
- The Remote Volume Mirroring feature is enabled, but the current number of Remote Volume Mirrors created on the primary storage array exceeds the number of mirror relationships allowed for this storage array.

Resolution:

- If the Remote Volume Mirroring feature is not enabled for the primary storage array:

1. Obtain a feature key for this feature from your storage supplier.
2. Enable the Remote Volume Mirroring premium feature. For step-by-step instructions, see “Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard” on page 143 or the **Array Management** window online help.
 - If the Remote Volume Mirroring feature is enabled, but the current number of mirror relationships on the storage array exceeds the number of mirror relationships allowed for the storage array, reduce the number of mirror relationships to equal or less than the number of mirror relationships allowed. There can be up to 32 defined mirrors on an SGI TP9300, up to 64 defined mirrors on an SGI TP9500./IS4000/IS350 storage array, and up to 128 defined mirrors on an SGI TP9700/IS4500 storage array.

For step-by step instructions, see “Removing Mirror Relationships” on page 154 or the **Array Management** window online help.

- **The remote volume candidate is unresponsive and cannot be updated.**

Cause:

The primary storage array is unable to notify the secondary storage array that its world wide name (WWN) has changed.

This may be the result of a controller in the primary storage array being swapped or removed. Storage array recovery procedures that result in configurations being reset could also cause a change to the WWN.

The first write to the primary volume will cause the mirrored volume pair to transition into an Unsynchronized status.

Resolution:

1. Remove all the existing mirror relationships from the primary storage array. For step-by-step instructions, see “Removing Mirror Relationships” on page 154 or the **Array Management** window online help.
2. Re-establish the required mirror relationships. For step-by step instructions, see “Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard” on page 143 or the **Array Management** window online help.

- **A long boot time exists on a Windows 2000 host that has secondary volumes mapped.**

Cause:

A Windows 2000 host can take up to two hours to boot if all of these statements are true:

- The storage management software is installed on the host.
- The Remote Mirroring feature is enabled.
- One or more secondary volumes are mapped to the host.

Resolution:

Do not map the secondary volumes to the host until they are needed for data access. If the host must be restarted, and the secondary volumes are mapped to that host, unmap the secondary volumes until the host has completed the restart. Remap the secondary volumes only when they are needed for data access.

Inappropriate Fibre Channel Fabric Configurations

This subsection describes Fibre Channel fabric configuration problems that may be encountered after you have configured your storage arrays and have activated and enabled the Remote Volume Mirroring feature.

Information about the probable cause and recommended resolution for each problem is provided. If your problem is not addressed in this section, see the remaining subsections in this section before contacting technical support.

- **The primary volume is unable to contact the Fibre Channel fabric.**

Cause:

Typically, this error occurs because of a failed Fibre Channel fabric or because of cabling between the primary and secondary storage arrays participating in a Remote Volume Mirror.

Resolution:

- If the Fibre Channel fabric has failed:
- Repair or replace the failed fabric.
- If the Fibre Channel fabric is operational:

Check the configurations cabling. The cable should not go through a hub or be connected directly between storage arrays. Ensure that the cable is attached through a switch.

For step-by-step instructions on configuring storage arrays for Remote Volume Mirroring, see “Hardware and Software Installation” on page 108.

- **The primary volume is unable to communicate with the secondary storage array. The primary storage array can still communicate with the fabric.**

Cause:

Typically, this error occurs when the secondary storage array has lost power or when network problems exist between the secondary storage array and the fabric.

Resolution:

- Secondary storage array has no power:
 - If the secondary storage array has no power and is not operational, restore power to the storage array. For information on powering on storage arrays, see your hardware documentation.
 - Secondary storage array has power and is operational:
 - Check the network status. Ensure the secondary storage array is visible on the network. Ensure the cables are correctly connected. For step-by-step instructions on configuring storage arrays for Remote Volume Mirroring, see “Hardware and Software Installation” on page 108.

Storage Array Component Failures

This subsection describes the impact of physical or logical storage array component failures on Remote Volume Mirroring functionality.

Information about the probable cause and recommended resolution for each problem is provided. If your problem is not addressed in this section, see the remaining subsections in this section before contacting technical support.

- **Following a role reversal, a snapshot volume created for a primary volume indicates that it has failed.**

Cause:

The base volume for the snapshot volume is a primary volume involved in a mirror relationship. A role reversal was performed, demoting the primary volume to secondary volume status, resulting in a snapshot volume on the secondary volume.

The base volume for a snapshot volume cannot be a candidate for the secondary volume role in a Remote Volume Mirror, and the snapshot volume is therefore failed.

Resolution:

Perform a role reversal to promote the secondary volume to the primary volume role. After the role reversal is complete, the snapshot volume will return to Optimal status.

For step-by-step instructions, see “Changing a Secondary Volume to a Primary Volume” on page 165 or the **Array Management** window online help.

- **The primary storage array controller indicates that it has failed.**

Cause:

Failure of the owning controller on the primary storage array will typically cause the associated multi-path driver on the attached host system to transfer ownership of the alternate controller on the primary storage array.

When this occurs, the alternate controller on the primary storage array will begin forwarding newly-written data to its peer on the secondary storage array. The secondary controller will take ownership of the volume from its (secondary) peer controller. This is a reactive event, which occurs because of the arrival of a write request from the primary storage array to the non-owning controller on the secondary storage array.

Resolution:

Allow the ownership transfer to complete. Operations will continue in a normal fashion, with the alternate (typically the non-preferred) controller on the primary storage array interacting with the alternate controller on the secondary storage array.

Failure of the non-owning controller on the primary storage array will have no direct impact on the status of any Remote Volume Mirrors.

- **The secondary storage array controller indicates that it has failed.**

Cause:

One of the following scenarios has occurred:

- Controller A on the secondary storage array has failed, while Controller A is the owner of the primary volume.
- Controller B on the secondary storage array has failed while Controller B is the owner of the primary volume.

In both scenarios, the mirrored volume pair transitions to the Unsynchronized status.

Resolution:

Restore the failed controller. For detailed recovery procedures, see the **Recovery Guru**.

Upon restoration of the failed controller, the primary storage array will automatically transition to a Synchronizing status to restore coherency between the volumes.

General Switch Troubleshooting

This subsection describes some common switch-related problems that may be encountered when using the Remote Volume Mirroring feature. The troubleshooting information provided relates to switches supported for use with this version of the storage management software. For a list of supported switches, see “Feature Concepts and Requirements” on page 79.

Information about the probable cause and recommended resolution for each problem is provided. If your problem is not addressed in this section, see the remaining subsections in this section before contacting technical support.

- **Host is unable to detect primary and secondary storage arrays participating in Remote Volume Mirroring.**

Cause:

When two or more switches are cascaded together and have the same Domain ID, hosts will be unable to locate the target device. Each switch must have a unique Domain ID when attempting to cascade two switches together.

Resolution:

Change the Domain ID of one or more of your switches, ensuring that each one is unique. The Domain ID should be a unique number between 1 and 239, which is used to identify the switch to a Fibre Channel fabric.

For more information on setting a Domain ID, see your switch user documentation.

Remote Volume Mirroring Critical Events

Critical Major Event Log (MEL) events are generated when an error occurs that can affect data availability or results in a degraded mode of operation. The storage management software responds by alerting the appropriate administrators using e-mail, SNMP trap, or other configured mechanisms.

Because administrative action is required to correct the problem, the storage array will generally enter a Needs Attention status. This causes the appropriate indications to be presented through the storage management software and an associated **Recovery Guru** procedure to be displayed when requested by the end user.

The following provides a description of critical MEL events, with a brief explanation of what caused the event to trigger, and the recommended resolution. If your problem is not addressed in this section, see the remaining sections in this chapter before contacting technical support.

- **Mirrored pair becomes unsynchronized.**

Cause:

This event is logged when a Remote Volume Mirror transitions to an Unsynchronized status, from either a Synchronization-in-Progress or Optimal status.

Resolution:

See the associated **Recovery Guru** message for specific recovery procedures.

- **Dual-role conflict detected.**

Cause:

This event is logged when a dual-role conflict is detected, and is typically the result of a role reversal. Since both sides of the mirrored volume pair are in the same primary or secondary volume role, both the primary and secondary storage arrays will report this MEL event.

Resolution:

See the associated **Recovery Guru** message for specific recovery procedures.

- **Communication error with the secondary volume on the secondary, remote storage array.**

Cause:

This event is logged when an interruption in the link that cannot be recovered by the firmware Fibre Channel driver is detected. Error codes and recovery procedures will assist the user in finding the exact cause of the communication problem between storage arrays participating in a mirror relationship.

Resolution:

See the associated **Recovery Guru** message for specific recovery procedures.

- **World wide name (WWN) change notification failed.**

Cause:

This event is logged when a storage array detects during start-up processing that its world wide name (WWN) has changed. The WWN can only change if both controllers of the array are exchanged while the array is powered off.

When the firmware detects this name change, it attempts to notify any remote storage array that had previously been participating in a mirroring relationship with any volume on the array.

Resolution:

See the associated **Recovery Guru** message for specific recovery procedures.

Upgrading Controller Firmware

The Remote Mirroring feature is designed so that it is not necessary for the primary storage system and the secondary storage system to be running the same level of controller firmware. All inter-controller operations (including those needed for propagation of newly written data from the primary volume to the secondary volume) have been implemented to support mixed firmware levels to the greatest extent possible. This condition allows firmware upgrades to be performed on both the primary storage system and the secondary storage system without any loss of data.

You receive error messages when certain configuration operations are not allowed between storage systems that have different firmware levels. Examples of the forbidden configuration options include establishing new mirrors and modifying mirror parameters. However, the key data availability operations, including promotion from the secondary volume to the primary volume, are allowed in such circumstances.

This section describes the procedure to upgrade the controller firmware (version 6.12 and later) and to make sure that the Remote Mirroring features work as intended. This is the preferred sequence for performing firmware upgrades on storage systems that participate in mirror relationships.

Note: Depending on your storage array platform, you may not be able to select controller firmware version 5.40 and later.

Note: By following this sequence, it is possible to avoid transitions of the mirrored volume pair to an Unsynchronized status, avoiding the need for a full resynchronization afterwards.

1. Stop I/O to applications that operate on all affected primary volumes temporarily inactive.
2. Download firmware to the secondary storage array.

3. Resume application access to primary volumes. There will be a period of time after this step where the primary storage array is operating with a different firmware level than the secondary storage array, but the inter-controller protocols for data replication will support such interactions.
4. Download firmware to the primary storage array.
5. Resume all I/O operations that were stopped in step 1.

For procedures on downloading the controller firmware, refer to the online help topics in the **Array Management Window**.

Chapter 4

The Volume Copy Premium Feature

This chapter includes the following sections:

- “Volume Copy Concepts” on page 183
- “Using the Volume Copy Premium Feature” on page 195
- “Volume Copy Scenarios” on page 218
- “Troubleshooting” on page 242
- “Glossary of Volume Copy Terms” on page 247

Volume Copy Concepts

This section introduces basic Volume Copy premium feature concepts, and describes how information is copied between volumes within a single storage array and how other premium features can be used in conjunction with the Volume Copy premium feature.

Introduction to the Volume Copy Premium Feature

The Volume Copy premium feature is used to copy data from one volume (the source) to another volume (the target) in a single storage array. This feature can be used to back up data, to copy data from volume groups that use smaller capacity drives to volume groups using greater capacity drives, or to restore snapshot volume data to the associated base volume.

This section introduces you to source volume and target volumes, and describes how they interact to copy data from one volume to another volume within a single storage array.

Note: Examples have been used throughout this chapter and the remainder of the guide to provide further explanation.

Source Volumes

When you create a volume copy, a copy pair is created, which consists of a source volume and a target volume that are located on the same storage array.

The source volume is the volume that accepts host I/O and stores data. The source volume can be a standard volume, snapshot volume, base volume of a snapshot volume, or a Remote Volume Mirror primary volume.

When a volume copy is started, data from the source volume is copied in its entirety to the target volume. The source volume is available for read I/O activity only while a volume copy has a status of In Progress, Pending, or Failed.

After the volume copy completes, the source volume becomes available to host applications for write requests. The target volume automatically becomes read-only to hosts, and write requests to the target volume will be rejected.

The Read-Only attribute can be changed after the volume copy is completed or has been stopped. For more information about the Read-Only attribute, refer to “Setting Target Volume Permissions” on page 216.

Target Volumes

Caution: A volume copy will overwrite data on the target volume. Ensure that you no longer need the data, or have backed up the data on the target volume, before starting a volume copy.

A target volume contains a copy of the data from the source volume. The target volume can be a standard volume, the base volume of a Failed or Disabled snapshot volume, or a Remote Volume Mirror primary volume.

When a volume copy is started, data from the source volume is copied in its entirety to the target volume.

While the volume copy has a status of In Progress, Pending, or Failed, read and write requests to the target volume will be rejected by the controllers. After the volume copy completes, the target volume automatically becomes read-only to hosts, and write requests to the target volume will be rejected.

The Read-Only attribute can be changed after the volume copy has completed or has been stopped. For more information about the Read-Only attribute, refer to “Setting Target Volume Permissions” on page 216.

Creating a Volume Copy

Creating a volume copy (copying data from the source volume to the target volume) is managed by the controllers and is transparent to host machines and applications. This section describes how data is copied using the storage management software, and the actions taken by the controllers during a volume copy.

The Create Copy Wizard guides you through the process of creating a new volume copy. The first dialog in the Wizard is the **Introduction** dialog, which allows you to select a source volume from the list of valid source volumes on the storage array.

The next dialog in the Create Copy Wizard is the **Select Target Volume and Copy Priority** dialog. It allows you to select the target volume for the volume copy and set the copy priority. The copy priority determines how much of the controller’s resources are used to complete the volume copy versus fulfill I/O activity.

The **Preview** dialog is the final dialog in the Wizard. It provides an overview of the selected source volume, target volume, and copy priority for the volume copy.

When you have completed the Wizard dialogs, the volume copy starts and the preferred controller of the source volume reads data from the source volume and writes it to the target volume, as shown in Figure 4-1 on page 186. **Operation in Progress** icons are displayed on the source volume and target volume while the volume copy is completing.

While a volume copy has a status of In Progress, Pending, or Failed:

- The source volume is available for read I/O activity only
- The target volume is not available for read or write I/O activity

After the volume copy has completed:

- The source volume is available for both read and write I/O activity
- The target volume is available for read I/O activity, unless the Read-Only attribute has been disabled. For more information about setting the target volume’s Read-Only attribute, refer to “Setting Target Volume Permissions” on page 216.

Note: A maximum of eight volume copies can be In Progress at one time. If more than eight volume copies are created, they will have a status of Pending until one of the volume copies with a status of In Progress completes.

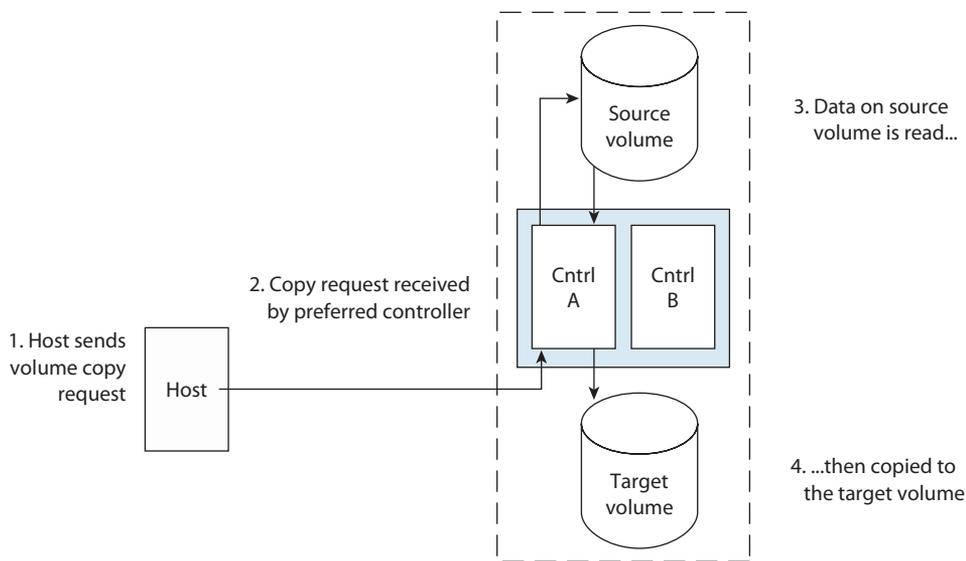


Figure 4-1 Creating a Volume Copy

While a volume copy is In Progress, the same controller must own both the source volume and target volume. If the same controller does not own both the source volume and target volume prior to the volume copy being created, ownership of the target volume is automatically transferred to the controller owner of the source volume.

When the volume copy is completed or is stopped, ownership of the target volume is restored to its preferred controller. If ownership of the source volume is changed during the volume copy, ownership of the target volume is also changed.

Volume Copy and Persistent Reservations

Note: Standard volumes holding a legacy or persistent reservation will not be available in the **Create Copy Wizard - Introduction** dialog or **Create Copy Wizard - Select Target** and **Set Copy Priority** dialog.

Persistent reservations are configured and managed through the server cluster software, and prevent other hosts from accessing the reserved volume.

Unlike other types of reservations, a persistent reservation reserves host access to the volume across multiple host ports, providing various levels of access control. The **Persistent Reservations** option enables you to query the storage array about registered host ports and reservations, and will retain reservation details in the event of a storage array power loss.

For more information about persistent reservations, refer to the **Array Management** window online help.

Storage Array Performance

During a volume copy, the storage array's resources may be diverted from processing I/O activity to completing a volume copy, which may affect the storage array's overall performance.

Several factors contribute to the storage array's performance, including I/O activity, volume RAID level, volume configuration (number of drives in the volume groups and cache parameters), and volume type (snapshot volumes may take more time to copy than standard volumes).

When you create a new volume copy, you will define the copy priority to determine how much controller processing time is allocated for a volume copy versus I/O activity.

There are five relative settings ranging from Lowest to Highest. The Highest priority rate supports the volume copy, but I/O activity may be affected. The Lowest priority rate supports I/O activity, but the volume copy will take longer.

The copy priority can be specified before the volume copy begins in the Select Target Volume and Copy Priority dialog of the Create Copy Wizard, as shown in Figure 4-2.

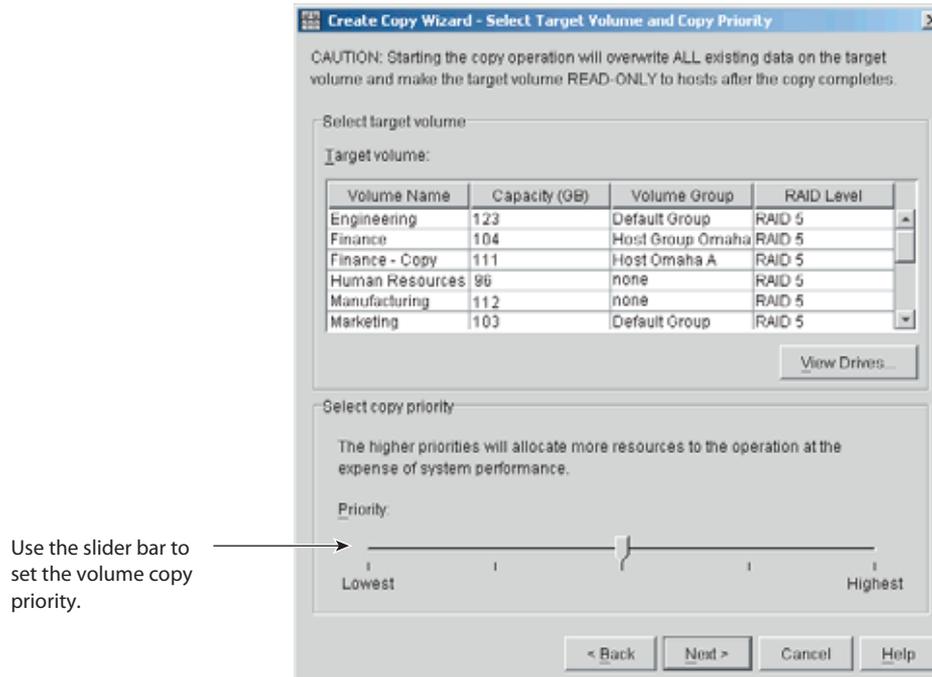


Figure 4-2 Select Target Volume and Copy Priority Dialog

While the volume copy is In Progress or after it has completed, the priority can be set in the **Change Copy Priority** dialog, as shown in Figure 4-3. The copy priority can also be specified in the **Re-Copy** dialog when a new volume copy for an existing copy pair is created.

For more information about the Re-Copy dialog, see “Re-Copying a Volume” on page 210.

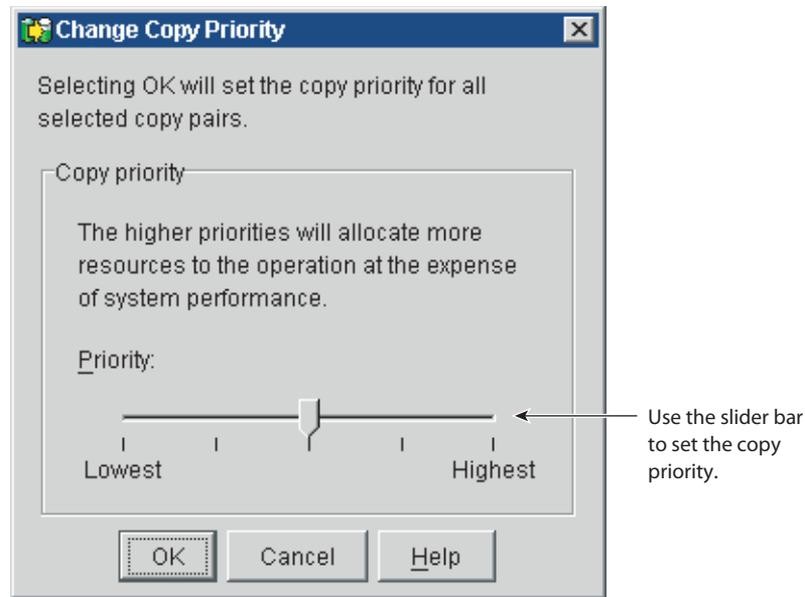


Figure 4-3 Change Copy Priority Dialog

Interaction of Volume Copy and Other Premium Features

This section describes how the Volume Copy premium feature can be used in conjunction with other premium features available in this version of the storage management software. Premium features covered in this section include Storage Partitioning, Snapshot, and Remote Volume Mirroring.

Storage Partitioning

Storage Partitioning is a premium feature that enables hosts to share access to volumes in a storage array. A storage partition is created when you define a collection of hosts (a host group) or a single host and then define a volume-to-LUN (Logical Unit Number) mapping.

This mapping allows you to define what host group or host will have access to a particular volume in the storage array.

Figure 4-4 shows the Mappings View of the **Array Management** window where the source volume **Finance** has been included in a partition accessible by **Host Omaha A**, and the target volume **Finance-Copy** has been included in a partition accessible by **Host Omaha B**.

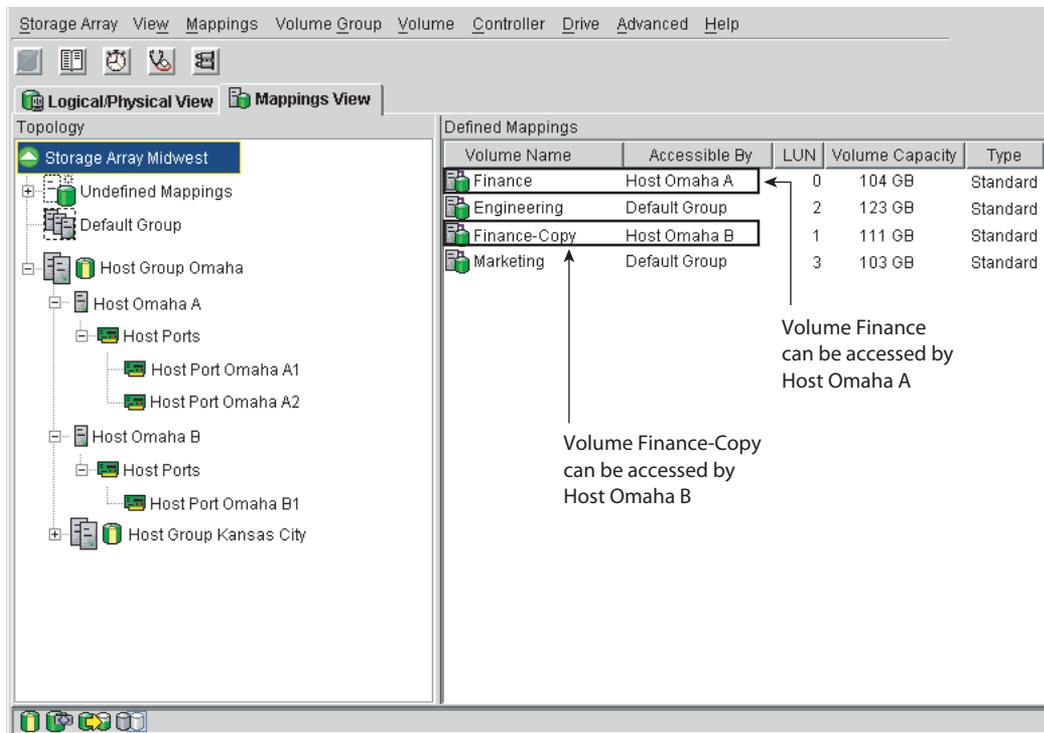


Figure 4-4 Storage Partitioning Example

After a volume copy has been created, the target volume automatically becomes read-only to hosts to ensure that the data is preserved. Hosts that have been mapped to a target volume will not have write access to the volume, and any attempt to write to the read-only target volume will result in a host I/O error.

If you want hosts to have write access to the data on the target volume, use the Copy Manager to disable the Read-Only attribute for the target volume. For more information about changing the target volume Read-Only attribute, see “Setting Target Volume Permissions” on page 216.

Snapshot Volume

A snapshot volume is a point-in-time image of a standard volume and is typically created so that an application, such as a backup application, can access the snapshot volume and read the data while the base volume remains online and accessible to hosts.

Snapshot volumes can be selected as the source volume for a volume copy. This is one of the best ways to use this feature, as it enables complete backups without any impact to the storage array I/O.

Note: If you choose the base volume of a snapshot volume as your target volume, you must disable all snapshot volumes associated with the base volume before you can select it as a target volume. For more information about disabling snapshot volumes, see “Volume Copy Scenarios” on page 218.

When you create a snapshot volume, a snapshot repository volume is automatically created. The snapshot repository volume stores information about the data that has changed since the snapshot volume was created. Snapshot repository volumes cannot be selected as a source volume or target volume in a volume copy.

The Snapshot Volume premium feature can be used in conjunction with the Volume Copy premium feature to back up data on the same storage array, and to restore the data on the snapshot volume back to its original base volume. For more information about volume copy usage scenarios, see “Volume Copy Scenarios” on page 218.

Remote Volume Mirroring

The Remote Volume Mirroring premium feature is used for online, real-time replication of data between storage arrays over a remote distance. In the event of a disaster or a catastrophic failure of one storage array, you can promote a second storage array to take over responsibility for computing services.

When you create a Remote Volume Mirror, a mirrored volume pair is created, consisting of a primary volume at the primary storage array and a secondary volume at a secondary, remotely located storage array.

The primary volume is the volume that accepts host I/O and stores data. When the mirror relationship is initially created, data from the primary volume is copied in its entirety to the secondary volume. This process is known as a full synchronization and is

directed by the controller owner of the primary volume. During a full synchronization, the primary volume remains fully accessible for all normal I/O activity.

The controller owner of the primary volume initiates remote writes to the secondary volume to keep the data on the two volumes synchronized. Whenever the data on the primary volume and the secondary volume becomes unsynchronized, the controller owner of the primary volume initiates a full synchronization.

The secondary volume maintains a mirror of the data on its associated primary volume. The controller owner of the secondary volume receives remote writes from the primary volume controller owner, but will not accept host read or write requests.

The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure at the primary site, you can perform a role reversal to promote the secondary volume to a primary role. Hosts will then be able to access the newly promoted volume and business operations can continue. For more information about performing a role reversals, see “Role Reversals” on page 193.

A primary volume participating in a Remote Volume Mirror can be selected as the source volume for a volume copy, as shown in Figure 4-5. A secondary volume participating in a Remote Volume Mirror cannot be selected a source volume or target volume and will not be available in the Create Copy Wizard when you create a volume copy.

For more information about the Remote Volume Mirroring premium feature, refer to the **Array Management** window online help.

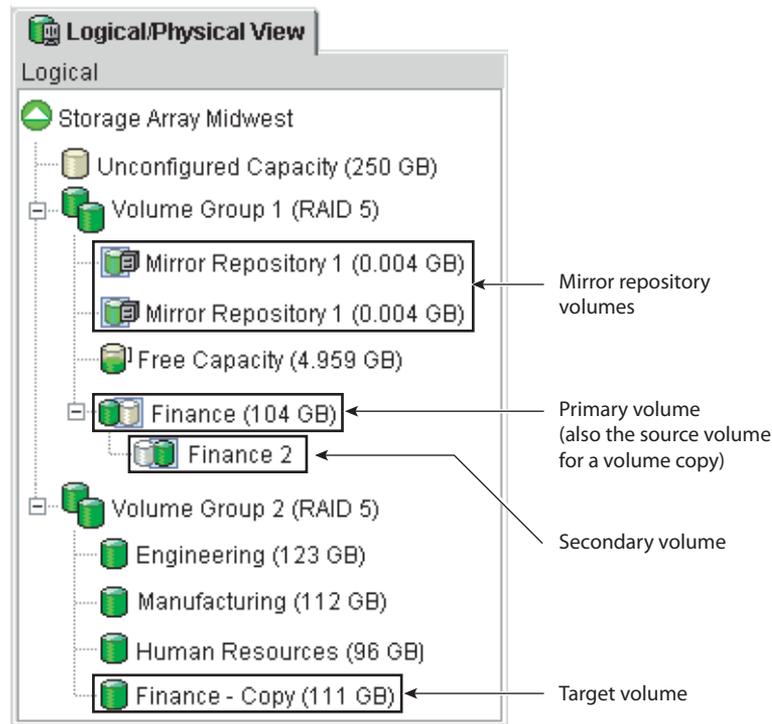


Figure 4-5 Primary Volume Included in a Volume Copy

Role Reversals

A role reversal is the act of promoting the secondary volume to be the primary volume of the mirrored volume pair, and demoting the primary volume to be the secondary volume.

In the event of a disaster at the storage array containing the primary volume, you can fail over to the secondary site by performing a role reversal to promote the secondary volume to the primary volume role. This will allow hosts to continue accessing data and business operations to continue.

A role reversal may be performed using one of the following methods:

- Change a secondary volume to a primary volume by selecting the **Volume >> Remote Volume Mirroring >>Change >> Role** to Primary option.

This option promotes the selected secondary volume to become the primary volume of the mirrored pair, and would be used when a catastrophic failure has occurred. For step-by-step instructions, refer to the **Array Management** window online help.

- Change a primary volume to a secondary volume by selecting the **Volume >> Remote Volume Mirroring >> Change >> Role** to Secondary option.

This option demotes the selected primary volume to become the secondary volume, and would be used during normal operating conditions. For step-by-step instructions, refer to the **Array Management** window online help.

If the primary volume is selected as the source volume for a volume copy, any volume copies that have a status of In Progress or Pending when the role reversal takes place will fail and cannot be started again until the primary volume is demoted back to its original Remote Volume Mirror secondary volume role.

If the primary storage array is recovered but is unreachable due to a link failure, a forced promotion of the secondary volume will result in both the primary and secondary volumes viewing themselves in the primary volume role (dual-primary condition). If this occurs, the volume copy in which the primary volume is participating will be unaffected by the role change.

Volume Copy Restrictions

The following restrictions apply to the source volume, target volume, and storage array:

- While a volume copy has a status of In Progress, Pending, or Failed, the source volume will be available for read I/O activity only. After the volume copy completes, read and write I/O activity to the source volume will be permitted.
- A volume can be selected as a target volume for only one volume copy at a time.
- The maximum allowable number of volume copies per storage array is dependent upon your storage array's configuration. Contact technical support for more information.
- A volume that is reserved by the host cannot be selected as a source volume or target volume.
- A volume with a status of Failed cannot be used as a source volume or target volume.
- A volume with a status of Degraded cannot be used as a target volume.

- A volume participating in a modification operation cannot be selected as a source volume or target volume. Modification operations include Dynamic Capacity Expansion (DCC), Dynamic RAID Level Migration (DRM), Dynamic Segment Sizing (DSS), Dynamic Volume Expansion (DVE), and defragmenting a volume.

Using the Volume Copy Premium Feature

This section provides step-by-step procedures for enabling the Volume Copy premium feature, creating volume copies, viewing volume copy properties, and managing volume copies using the Copy Manager.

Enabling the Volume Copy Premium Feature

Before you can begin creating volume copies, you will need to enable the Volume Copy premium feature. To complete this procedure, you will need to obtain a Feature Key File from your storage supplier.

If the Volume Copy premium feature is disabled, you cannot create new volume copies.

If the Volume Copy premium feature is disabled, but volume copy pairs exist, all of these options in Copy Manager are still available:

- Stop copy
 - Remove copy pairs
 - Start copy
 - Enable read-only
 - Disable read-only
1. From the Array Management window, select **Storage Array >> Premium Features >> Enable**.

The **Select Feature Key File** dialog is displayed.

Note: The **Select Feature Key File** dialog filters for files with the *.key extension.

2. Select the folder in which you placed the feature key file.

3. Select the appropriate feature key file, and then select **OK**.

The **Enable Premium Feature** dialog box is displayed.

4. Select **Yes**.

The Volume Copy premium feature is now enabled. The icon in the premium feature status area is no longer displayed with a red slash. To further verify the status of the feature, select **Storage Array >> Premium Features >> List**.

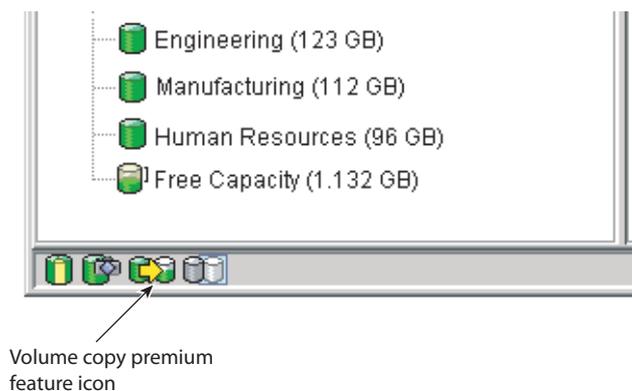


Figure 4-6 Enabling the Volume Copy Premium Feature

If the Volume Copy premium feature is disabled, but volume copy pairs exist, all of the options in the Copy Manager (stop copy, remove copy pairs, start copy, enable and disable read-only) will still be available. However, you will not be permitted to create any new volume copies.

Creating a Volume Copy

This section introduces you to selecting or creating a target volume, selecting the source volume, creating the volume copy, and managing the volume copy using the Copy Manager.

Selecting or Creating a Target Volume

Caution: A volume copy will overwrite data on the target volume. Ensure that you no longer need the data, or have backed up the data on the target volume, before starting a volume copy.

Before creating a volume copy, ensure that a suitable target volume exists on the storage array, or create a new target volume specifically for the volume copy. Only valid target volumes with a capacity equal to or greater than the source volume will be displayed in the **Create Copy Wizard - Select Target Volume and Copy Priority** dialog, in alphanumeric order.

The following are valid target volumes and will be displayed in the **Create Copy Wizard - Select Target Volume and Copy Priority** dialog:

- Standard volume
- Base volume of a disabled or failed snapshot volume
- Primary volume participating in a Remote Volume Mirror

The following are not valid target volumes and will not be displayed in the **Create Copy Wizard - Select Target Volume and Copy Priority** dialog:

- Base volume of an active snapshot volume
- Snapshot volume
- Mirror repository volume
- Snapshot repository volume
- Secondary volume in a Remote Volume Mirror
- Failed volume
- Missing volume

- A volume with a status of Degraded
- A volume currently in a modification operation
- A volume holding a legacy or persistent reservation
- A volume that is a source volume or target volume in another volume copy that has a status of In Progress, Pending, or Failed

To make locating the target volume in the Logical View easier, ensure that you add “target” or “copy” to the target volume’s name, as shown in Figure 4-7. For example, if the source volume name is Finance, name the target volume Finance-Copy. That way you can quickly identify the source volumes and target volumes available on the storage array.

Once the volume copy has completed, the target volume automatically becomes read-only to hosts. Any write requests to the target volume will be rejected, unless you disable the Read-Only attribute in the Copy Manager. For more information about disabling the target volume’s Read-Only attribute, see “Setting Target Volume Permissions” on page 216.

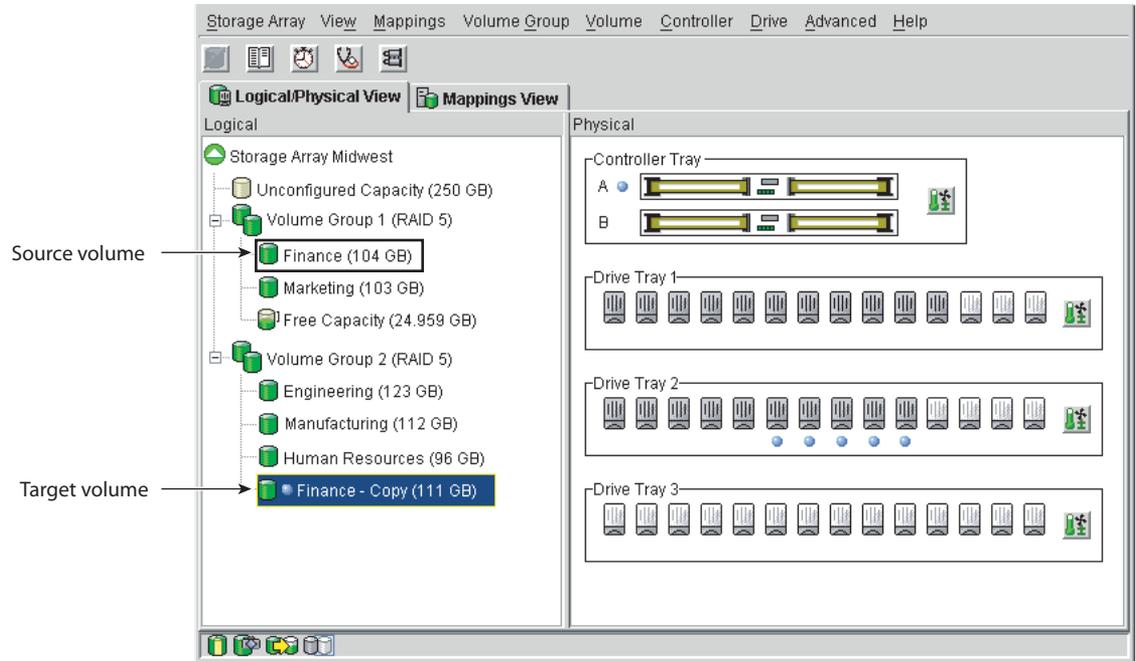


Figure 4-7 Renaming the Target Volume

Selecting a Source Volume

The first dialog of the Create Copy Wizard allows you to choose the source volume from which data will be copied. A source volume is a standard volume that accepts host I/O and stores data.

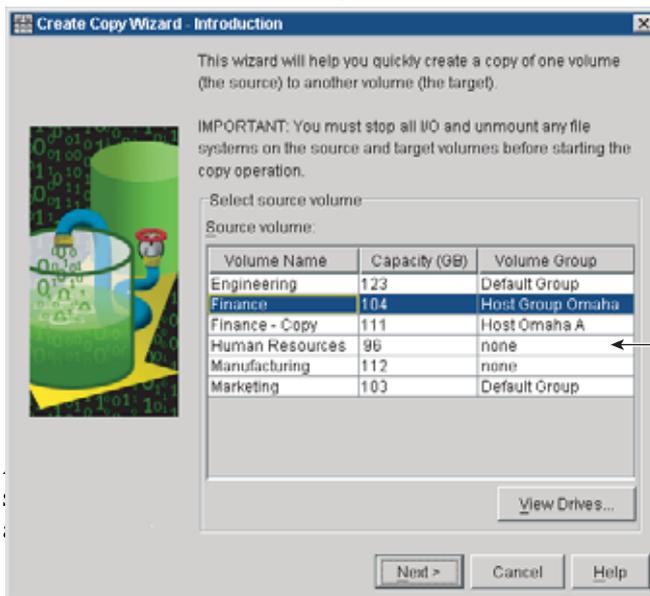
Note: If the volume you selected is not a valid source volume, the Create Copy Wizard will not be displayed. You will need to select a valid volume to begin the Wizard.

When the volume copy is started, data from the source volume is copied in its entirety to the target volume. While a volume copy has a status of In Progress, Pending, or Failed, the source volume will be available for read I/O activity only. After the volume copy completes or is stopped, read and write access to the source volume will be permitted.

The following are valid source volumes and will be displayed in the **Create Copy Wizard - Introduction** dialog:

- Standard volume
- Snapshot volume
- Base volume of a snapshot volume
- Primary volume participating in a Remote Volume Mirror

The following are not valid source volumes and will not be displayed in the **Create Copy Wizard - Introduction** dialog:



Valid source volumes are listed in order of increasing capacity.

mes on the me capacity,

Figure 4-8 Selecting a Source Volume

Viewing Drives

The **View Drives** dialog displays the drive tray, slot, and capacity information for the selected source volume or target volume. The **View Drives** dialog is available in the **Create Copy Wizard - Introduction** and **Create Copy Wizard - Select Target Volume** dialogs.

Content for this dialog will automatically update dynamically if you select a different source volume or target volume in the **Copy Wizard** dialog while the **View Drives** dialog is open.

The information presented in the **View Drives** dialog can be used to copy high-access volumes to different volume groups for increased accessibility, or to copy volumes to a volume group that uses drives with a greater capacity.

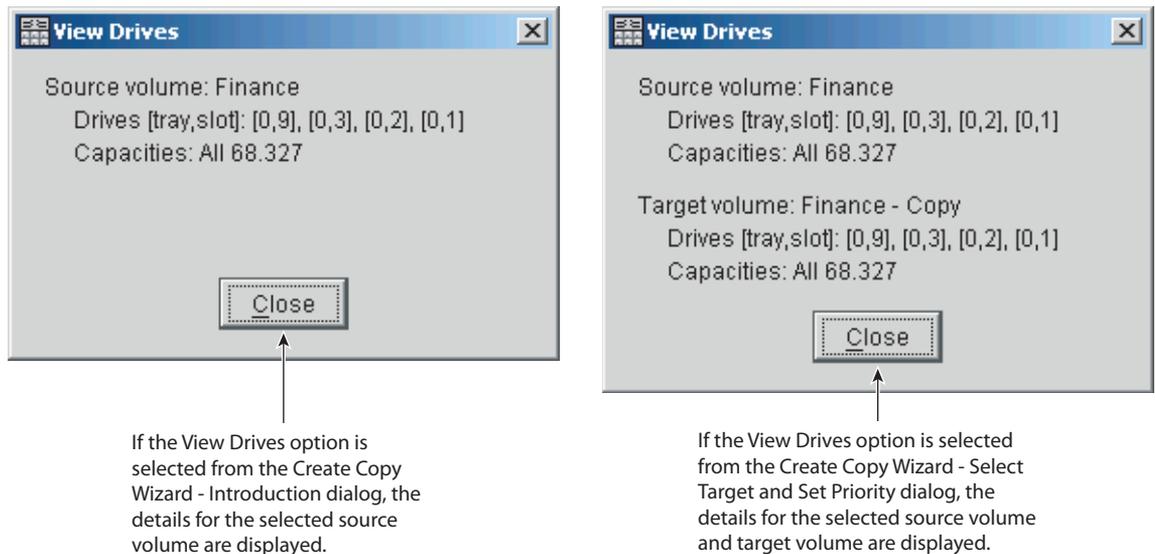


Figure 4-9 Examples of the View Drives Dialog

Creating a Volume Copy Using the Create Copy Wizard

The Create Copy Wizard allows you to create a full copy of a source volume, overwriting any existing data on the target volume. To create a volume copy, complete the following steps.

1. Stop all I/O activity to the source volume and target volume.
2. Unmount any filesystems on the source volume and target volume, if applicable.
3. In the **Logical View** of the **Array Management** window, select the volume that will be the source volume for the volume copy.
4. **Select Volume >> Copy >> Create** from the menu.

The **Create Copy Wizard - Introduction** dialog is displayed, as shown in Figure 4-10.

Note: If the volume you selected is not a valid source volume, the Create Copy Wizard will not be displayed. You will need to select a valid volume to begin the Wizard.

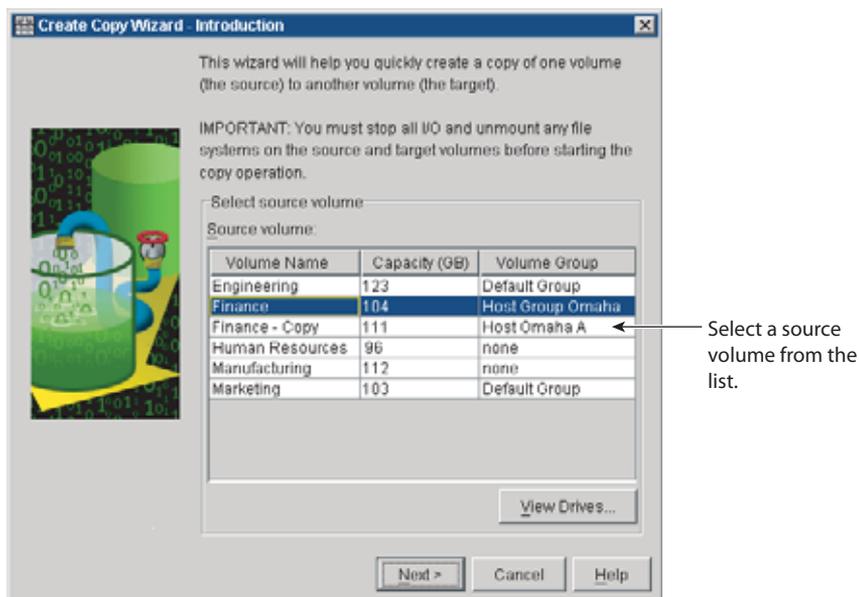
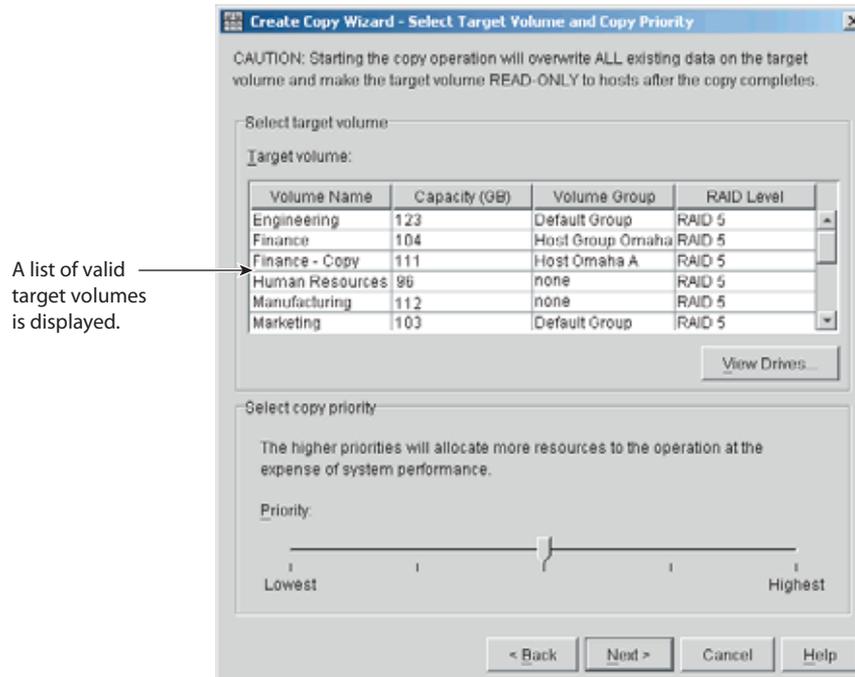


Figure 4-10 Create Copy Wizard - Introduction Dialog

5. **Select Next.**

The **Create Copy Wizard - Select Target Volume and Copy Priority** dialog is displayed, as shown in Figure 4-11.

Caution: A volume copy will overwrite data on the target volume. Ensure that you no longer need the data, or have backed up the data on the target volume, before starting a volume copy.



A list of valid target volumes is displayed.

Figure 4-11 Create Copy - Select Target Volume and Copy Priority Dialog

6. Select the target volume from the list of available volumes.

Only valid target volumes are listed, with the volume closest in size to the source volume highlighted.

Note: If you select a target volume with a capacity similar to the source volume, you will reduce the risk of having unusable space on the target volume after the volume copy completes.

7. Set the copy priority for the volume copy.

The higher priorities will allocate more resources to the volume copy at the expense of the storage array's performance.

8. Select **Next**.

The **Create Copy Wizard - Preview** dialog is displayed, as shown in Figure 4-12. The source volume and target volume you selected are displayed, together with the selected copy priority.

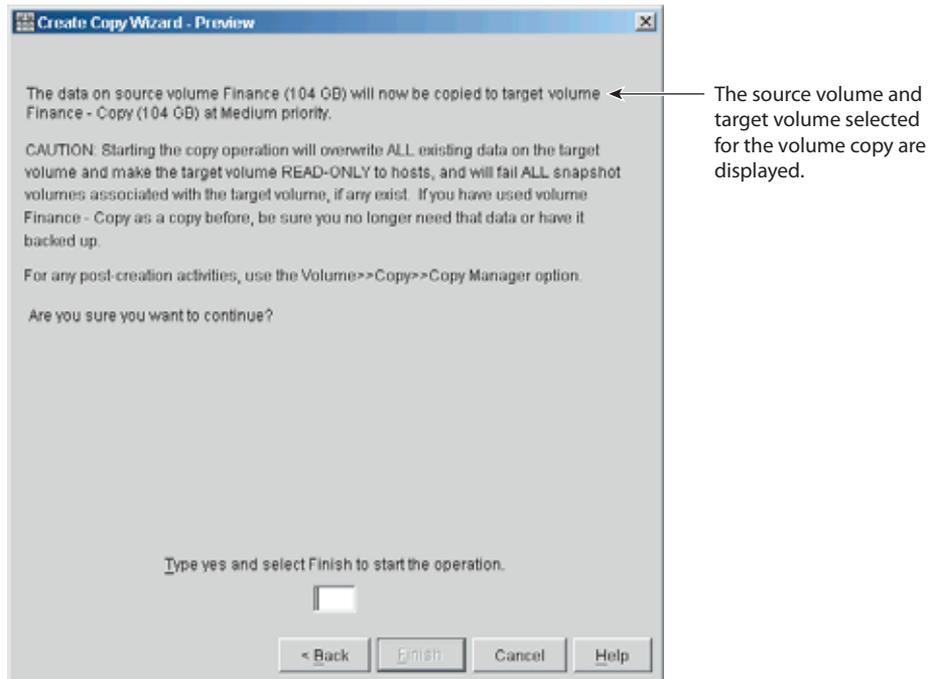


Figure 4-12 Create Copy Wizard - Preview Dialog

9. Review the information displayed in the **Preview** dialog. Type **Yes** and select **Finish** to start the volume copy.

The **Copy Started** dialog is displayed, verifying that the volume copy has started, as shown in Figure 4-13. This dialog also gives you the option to create a new volume copy or exit the Wizard.

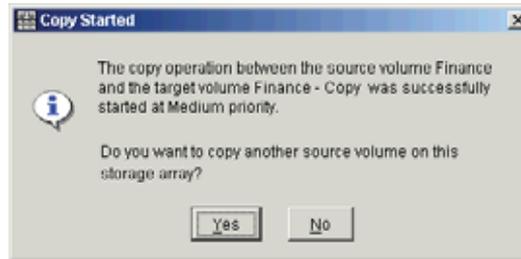


Figure 4-13 Copy Started Dialog

10. Choose one of the following, based on whether you want to create another volume copy:
 - **Yes** – create a new volume copy.
 - **No** – exit the Wizard and return the **Logical/Physical** view.
11. Remount any filesystems on the source volume and target volume.

SMrepassist Utility for Windows Operating Systems

SMrepassist is a host-based utility for Windows operating systems (OSs). Use the SMrepassist utility before and after you create a volume copy on the Windows OS. The SMrepassist utility makes sure that all of the memory-resident data for the file systems on the target volume is flushed. This utility also makes sure that the driver recognizes the signatures and the file system partitions.

You can use the SMrepassist utility to resolve duplicate signature problems for snapshot volumes. To run the SMrepassist utility, type the following command, and press Enter:

```
SMrepassist -f [filesystem-identifier]
```

- where `-f` flushes all the memory-resident data for the file system that is indicated by [*filesystem-identifier*]
- where [*filesystem-identifier*] specifies a unique file system by using this syntax:
drive-letter: [*mount-point-path*]

Note: The file system can include only a drive letter, or a drive letter and a mount point.

When the `SMrepassist` utility cannot differentiate between the following items, an error message appears on the command line, and the volumes are listed as a volume set:

- Base volume and snapshot volume (for example, if the snapshot volume was removed)
- Standard volume and volume copy (for example, if the volume copy was removed)

Locating Associated Volumes

The Go To option can be used to quickly jump to an associated source volume or target volume displayed in the **Logical View** of the **Array Management** window.

Note: This option is available only if the Volume Copy premium feature has been enabled, or if volume copies currently exist on the storage array. It is not accessible using the **Mappings View** of the **Array Management** window.

To locate an associated source volume or target volume, complete the following steps.

1. Select a source volume or target volume in the **Logical View** of the **Array Management** window.
2. Choose one of the following, based on the volume type selected:
 - **Source volume** – Select the **View >> Go To >> Target Volume** menu option
The selection jumps to the associated target volume in the Logical View of the Array Management Window
 - **Target volume** – Select the **View >> Go To >> Source Volume** menu option
The selection jumps to the associated source volume in the **Logical View** of the **Array Management** window.

Viewing Volume Copy Properties

The **Volume Properties** dialog enables you to view information about a selected source volume or target volume.

This dialog displays the role, copy status, copy priority, start timestamp, completion timestamp, the target volume's read-only attribute setting, the source volume or target volume World Wide Name (WWN), and if the volume copy is in progress, a progress bar with the estimated completion time for the volume copy is displayed.

To view the volume copy properties for a source volume or target volume, complete the following steps.

1. Select the source volume or target volume in the **Logical View** of the **Array Management** window.
2. Select **Volume >> Properties** to display the **Volume Properties** dialog, as shown in Figure 4-14.

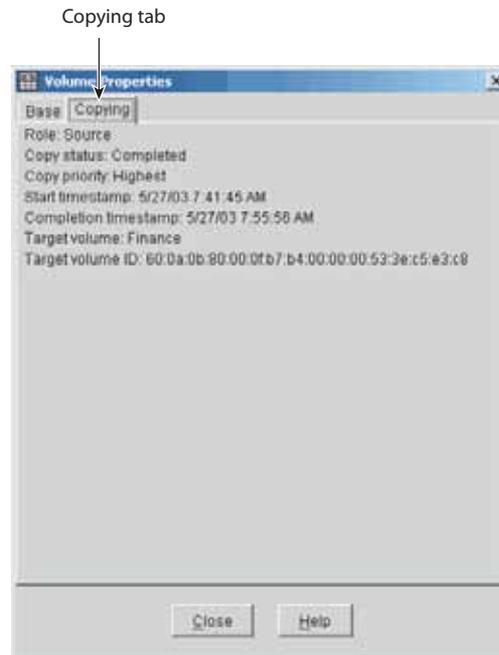


Figure 4-14 Volume Properties Dialog

3. Select the **Copying** tab.

The **Copy Properties** tab is displayed. The volume properties displayed include:

- Role: Source or Target
- Copy Status: Pending, In Progress, Completed, Stopped, Failed
- Copy Priority: Lowest, Low, Medium, High, Highest
- Start Timestamp: MM/DD/YY HH/MM/SS
- Completion Timestamp: MM/DD/YY HH/MM/SS
- Read-Only: Enabled/Disabled (only if target volume is selected)
- Source/Target Volume: Volume Name
- Source/Target Volume ID: WWN (World Wide Name)

4. Select **Close** to exit the dialog.

If a volume is participating in more than one volume copy (can be a source volume for one volume copy and the target volume for another volume copy), the details will be repeated for each associated copy pair.

Managing Volume Copies with Copy Manager

The **Copy Manager** dialog, launched from the **Array Management** window, monitors the progress of a volume copy and allows you to manage volume copies while they are in progress and after they have been created.

Information is displayed for all the volume copies on the storage array, including the source volume from which data is being copied, the target volume for the copied data, the status of the volume copy, a timestamp for any completed volume copies, and the volume copy priority that was assigned to the volume copy, as shown in Figure 4-15.

The source volume and target volume for a single volume copy are known as a *copy pair*.

Note: A maximum of eight volume copies can have a status of In Progress at one time. If more than eight volume copies are created, they will each have a status of Pending until one of the volume copies with a status of In Progress completes.

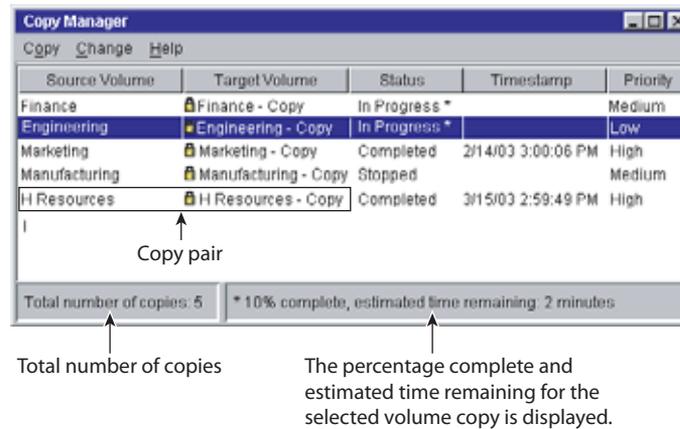


Figure 4-15 Copy Manager Dialog

The progress of the volume copy is displayed in the status bar at the bottom of the **Copy Manager** dialog as a percentage, with an estimation of the time remaining for the volume copy to complete. The total number of volume copies on the storage array is also displayed in the status bar at the bottom of the **Copy Manager** dialog.

The Copy Manager can be used to re-copy a volume copy, stop a volume copy with a status of In Progress, remove copy pairs (which removes associated volume copy configuration information from the storage array), change the volume copy priority, and disable the target volume's Read-Only attribute.

To manage a volume copy using the Copy Manager, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**.
The **Copy Manager** dialog is displayed.
2. Select a copy pair in the table by clicking on a specific copy pair.
3. Select an option from the pull-down menu bar to manage the highlighted volume copy.

The Copy Manager automatically refreshes when the information for a volume copy changes, and will poll the storage array at regular intervals when a volume copy has a status of In Progress to display the most accurate information possible.

Re-Copying a Volume

Caution: This option will overwrite existing data on the target volume and make the target volume read-only to hosts. This option will fail all snapshot volumes associated with the target volume, if any exist.

The Re-Copy option enables you to create a new volume copy for a previously defined copy pair that may have been Stopped, Failed, or has Completed. This option can be used for creating scheduled, complete backups of the target volume that can then be copied to tape drive for off-site storage.

After starting the Re-Copy option, the data on the source volume is copied in its entirety to the target volume.

You can also set the copy priority for the volume copy at this time. The higher priorities will allocate storage array resources to the volume copy at the expense of the storage array's performance.

To create a new volume copy for an existing copy pair, complete the following steps.

1. Stop all I/O to the source volume and target volume.
2. Unmount any filesystems on the source volume and target volume, if applicable.
3. Select **Volume >> Copy >> Copy Manager**.

The **Copy Manager** dialog is displayed.

Note: Only one copy pair at a time can be selected to be re-copied.

4. Select the copy pair in the table by clicking on a specific copy pair.
5. Select **Copy >> Re-Copy** from the **Copy Manager** menu.

The **Re-Copy** dialog is displayed, as shown in Figure 4-16.

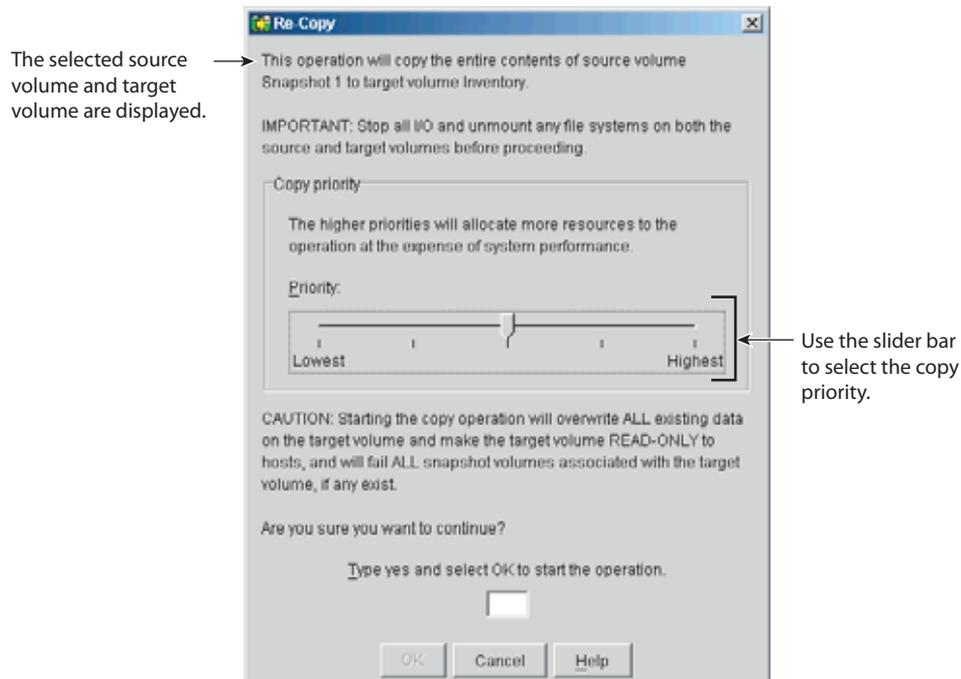


Figure 4-16 Re-Copy Dialog

6. Set the copy priority, as needed.

This setting defines how much of the system resources are used to complete a volume copy versus fulfill I/O requests. There are five relative settings ranging from Lowest to Highest. The Highest priority rate supports the volume copy, but I/O activity may be affected. The Lowest priority rate supports I/O activity, but the volume copy will take longer.

7. Type *Yes* and select **OK**.

The volume copy begins. The volume copy has a status of In Progress or Pending in the Copy Manager, as shown in Figure 4-17. In the **Logical View**, Operation in Progress icons are displayed on the source volume and target volume.

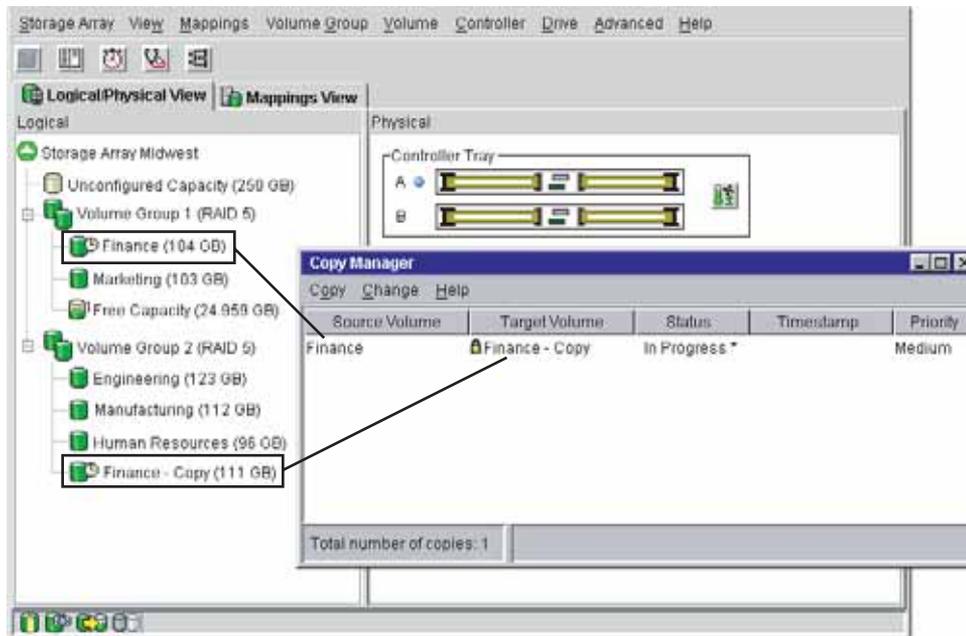


Figure 4-17 Re-Copying a Volume

Stopping a Volume Copy

The Stop Copy option allows you to stop a volume copy that has a status of In Progress, Pending, or Failed.

Using this option on a volume copy with a status of Failed clears the Needs Attention status displayed for the storage array in the **Logical View** of the **Array Management** window.

After the volume copy has been stopped, the Re-Copy option can be used to create a new volume copy using the original copy pair. For information about recreating a volume copy, see “Re-Copying a Volume” on page 210.

To stop a volume copy, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**, as shown in Figure 4-18.
The **Copy Manager** dialog is displayed.

2. Select the copy pair in the table by clicking on a specific copy pair.
3. Select **Copy >> Stop** from the **Copy Manager** menu.

The Stop Copy dialog is displayed.

Note: When the volume copy is stopped, all mapped hosts will have write access to the source volume. If data is written to the source volume, the data on the target volume will no longer match the data on the source volume.

4. Select **Yes** to stop the volume copy for the selected volumes.

The volume copy is displayed in the Copy Manager with a status of Stopped.

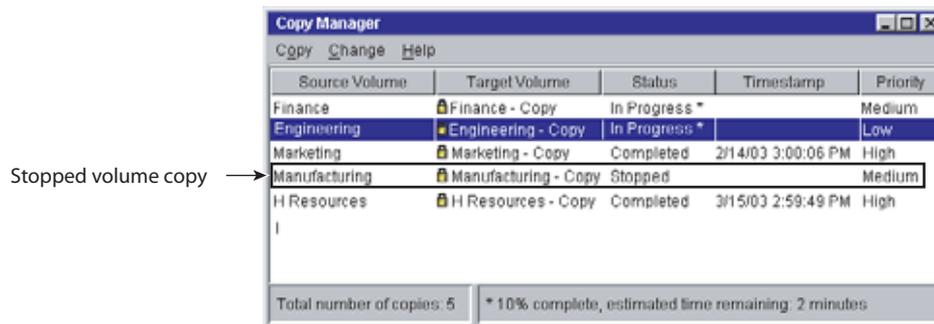


Figure 4-18 Stopping a Volume Copy

Removing Copy Pairs

The Remove Copy Pairs option allows you to remove one or more volume copies from the Copy Manager. Any volume copy-related information for the source volume and target volume is removed from the **Volume Properties** and **Storage Array Profile** dialogs.

After the volume copy is removed from the Copy Manager, the target volume can be selected as a source volume or target volume for a new volume copy.

Removing a volume copy from the Copy Manager also permanently removes the Read-Only attribute for the target volume.

Note: If the volume copy has a status of In Progress, you must stop the volume copy before you can remove the copy pair from the Copy Manager dialog.

To remove volume copies from the Copy Manager, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**.
The **Copy Manager** dialog is displayed.
2. Select one or more copy pairs in the table.
3. Select **Copy >> Remove Copy Pairs** from the **Copy Manager** menu.
The **Remove Copy Pairs** dialog is displayed.
4. Select **Yes** to permanently remove the copy pair from the Copy Manager.

Note: The data on the source volume or target volume is not deleted.

Changing Copy Priority

The **Change Copy Priority** dialog allows you to set the rate at which the volume copy completes, as shown in Figure 4-19. The copy priority setting defines how much of the storage array's resources are used to complete a volume copy versus fulfill I/O requests.

There are five relative settings ranging from Lowest to Highest. The Highest priority rate supports the volume copy, but I/O activity may be affected. The Lowest priority rate supports I/O activity, but the volume copy will take longer.

You can change the copy priority for a copy pair before the volume copy begins, while the volume copy has a status of In Progress, or after the volume copy has completed when recreating a volume copy using the Re-Copy option.

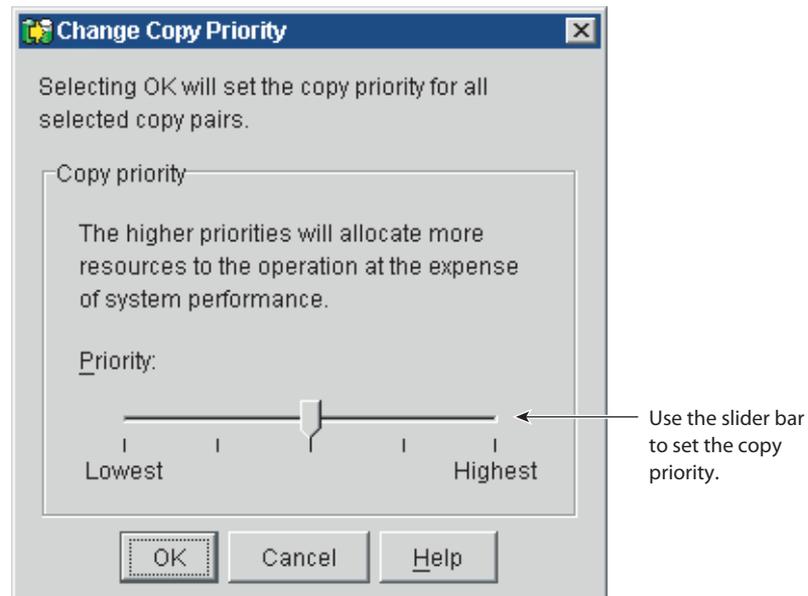


Figure 4-19 Change Copy Priority Dialog

To change the copy priority for a volume copy, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**.
The **Copy Manager** dialog is displayed.
2. Select the copy pair in the table by clicking on a specific copy pair.
3. Select **Change >> Copy Priority** from the **Copy Manager** menu.
The **Change Copy Priority** dialog is displayed.
4. In the **Copy Priority** section, use the slider bar to select the appropriate copy priority, as shown in Figure 4-19.
5. There are five choices: Lowest, Low, Medium, High, and Highest. These settings are shown with five marks on the slider bar. The Lowest and Highest settings are labeled. The higher the setting, the larger the impact to system performance.
6. Select **OK**.
The **Change Copy Priority - Progress** dialog is displayed.
7. Select **OK**.

The copy priority for the selected volume copy is changed.

Setting Target Volume Permissions

Read and write requests to the target volume will be rejected while the volume copy has a status of In Progress, Pending, or Failed.

After the volume copy has completed, the target volume automatically becomes read-only to hosts. You may want to keep the Read-Only attribute enabled in order to preserve the data on the target volume.

Examples of when you may want to keep the Read-Only attribute enabled include:

- If you are using the target volume for backup purposes
- If you are copying data from one volume group to a larger volume group for greater accessibility
- If you are planning to use the data on the target volume to copy back to the base volume in case of a disabled or failed snapshot volume

If you decide to allow host write access to the data on the target volume after the volume copy is completed, use the Copy Manager to disable the Read-Only attribute for the target volume.

Disabling Target Volume Read-Only

The Disable Read-Only option allows you to change the target volume permissions, allowing host read and write access to the target volume after the volume copy is completed.

To disable the Read-Only attribute for the target volume, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**.
The **Copy Manager** dialog is displayed.
2. Select the copy pair in the table by clicking on a specific copy pair.
3. Select **Change >> Target Volume Permissions >> Disable Read-Only** from the **Copy Manager** menu.

The Read-Only attribute for the target volume is disabled, and the lock icon displayed next to the target volume in the Copy Manager is removed, as shown in Figure 4-20. Write requests are allowed to the target volume.

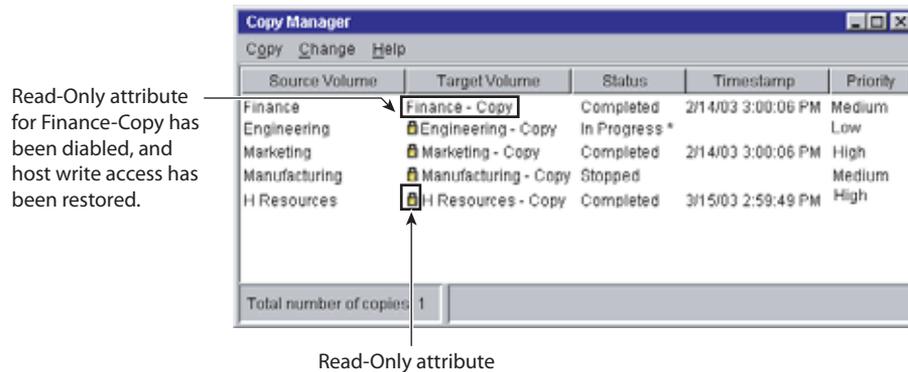


Figure 4-20 Disabling the Target Volume's Read-Only Attribute

Enabling Target Volume Read-Only

The Enable Read-Only option allows you to change the target volume permissions, preventing host write access to the target volume after the volume copy is completed, or prior to starting the volume copy.

To enable the Read-Only attribute for the target volume, complete the following steps.

1. Select **Volume >> Copy >> Copy Manager**.
The **Copy Manager** dialog is displayed.
2. Select the copy pair in the table by clicking on a specific copy pair.
3. Select **Change >> Target Volume Permissions >> Enable Read-Only** from the **Copy Manager** menu.

The Read-Only attribute for the target volume is enabled, and the lock icon is displayed next to the selected target volume. Write requests to the target volume are rejected.

Volume Copy Scenarios

This chapter describes three key usage scenarios for the Volume Copy premium feature and the major tasks associated with each scenario.

Using the Volume Copy premium feature, you will be able to relocate data for better performance, create complete copies for backups, and be able to restore snapshot volume data.

Copying Volumes for Improved Access

As your storage requirements for a volume change, the Volume Copy premium feature can be used to copy data to a volume in a volume group that uses greater capacity drives within the same storage array.

Copying data for greater access allows you to move data to drives:

- With higher capacity (73GB to 146 GB)
- With a higher data-transfer rate (1 Gb/s to 2 Gb/s)
- Using new technologies for higher performance

Before you select the target volume to which the data will be moved, make sure that the capacity of the target volume is greater than the capacity of the source volume.

In the following example, Storage Array Midwest consists of a command module and threedrive modules: Drive Modules 1 and 2 contain fourteen 73GB drives, and Drive Module 3 contains fourteen 146GB drives.

Volume Finance is associated with the drives in slot 3 on Drive Module 1 and slot 3 in Drive Module 2. To improve efficient access to the data on Finance, a volume copy will be created to copy the data to drives on Drive Module 3, as shown in Figure 4-21).

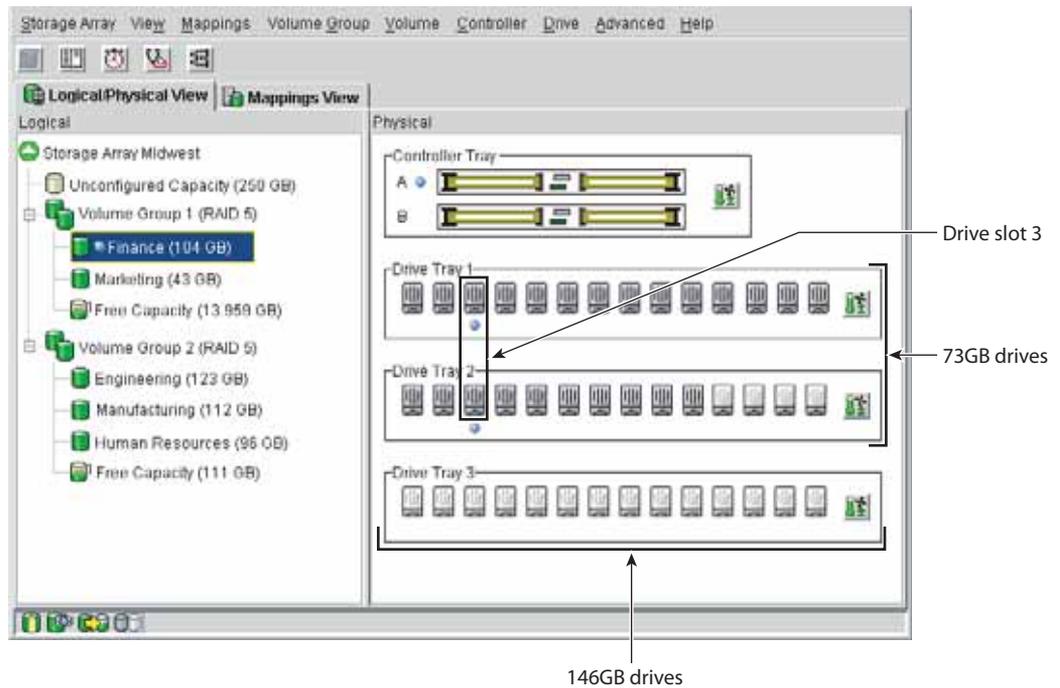


Figure 4-21 Storage Array Midwest Example

To copy data to a volume in a volume group using greater capacity drives, complete the following steps.

1. Use the Create Volume Wizard to create a new volume with the required parameters (RAID level, overall capacity, number of drives) on a volume group associated with the greater capacity drives, as shown in Figure 4-22. This will be the *target volume* for the volume copy.
 - a. Select the Free Capacity node on Volume Group 2.
 - b. Select **Volume >> Create** to display the Create Volume Wizard.
 - c. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

Note: To locate the target volume in the **Logical/Physical View** tab more easily, you can add the word “target” to the name of the target volume. For more information about how to create a standard volume, refer to the online help topics in the **Array Management Window**.

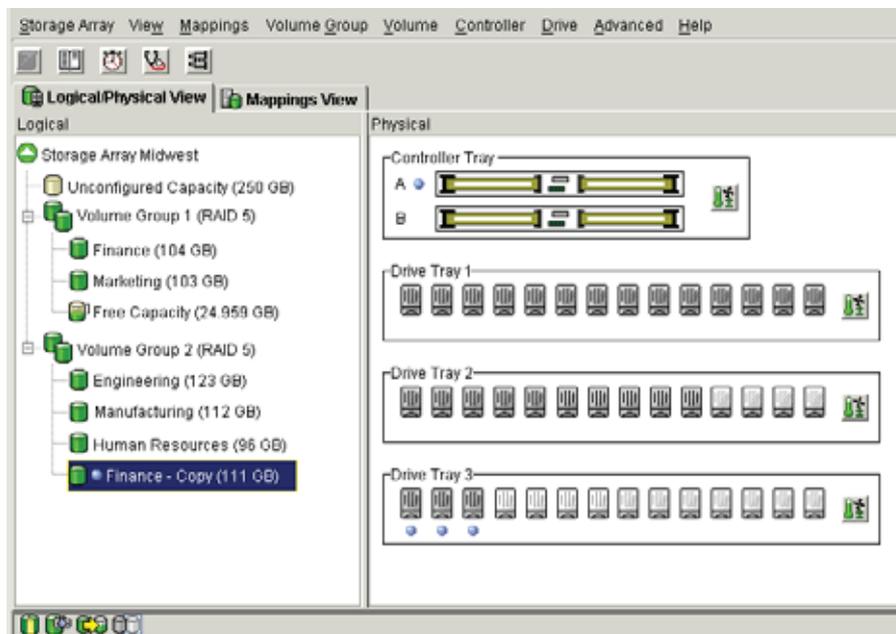


Figure 4-22 Creating the Target Volume Finance-Copy

2. Stop all I/O activity to the source volume.
3. Unmount any file systems on the source volume and target volume. Windows OS - use the SMrepassist utility to flush the cache of the source volume before you create a volume copy.
4. Delete the volume-to-LUN mapping for the source volume. Ensure that you record the LUN number used to provide access to the source volume. You will need this in step 8 when you define a volume-to-LUN mapping for the target volume.
 - a. From the Mappings View select the source volume, and then select **Mappings >> Delete** to display a confirmation dialog.

- b. Select **Yes** to delete the selected volume-to-LUN mapping and close the **Delete** confirmation dialog.

The volume-to-LUN mapping is deleted and the source volume is moved back under the Undefined Mappings node in the **Mappings View**.

5. Use the Create Copy Wizard to create a new volume copy, as shown in Figure 4-23.
 - c. Select **Volume >> Copy >> Create** to display the **Create Copy Wizard - Introduction** dialog. The source volume will be highlighted in the table.
 - d. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

Select the volume you created in step 1 as the target volume.

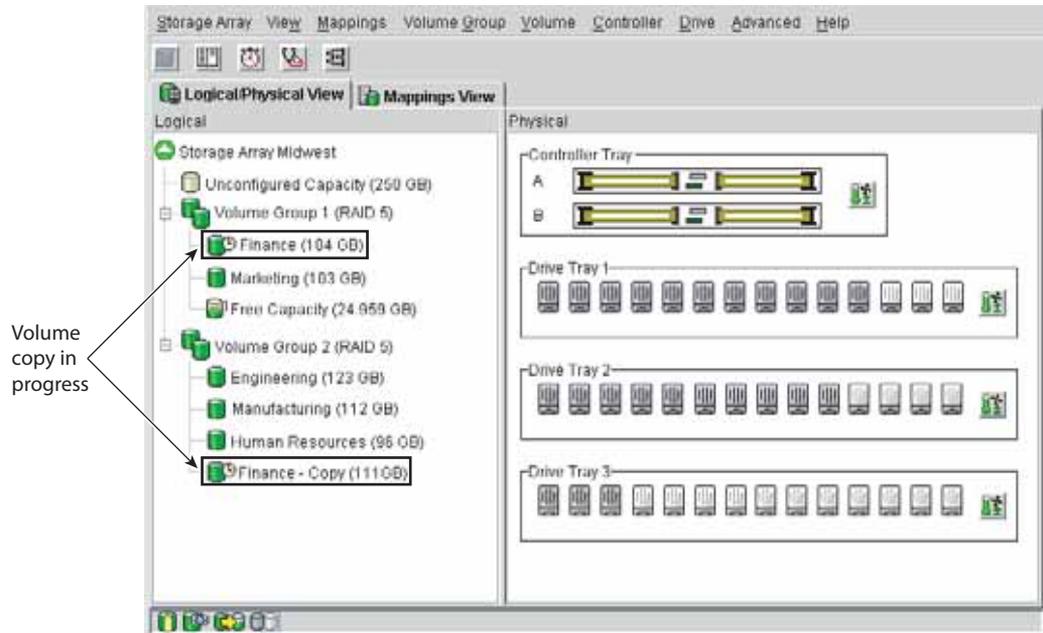


Figure 4-23 Creating a Volume Copy of Finance

6. Select **Volume >> Copy >> Copy Manager** to open the Copy Manager to view the status of the volume copy, as shown in Figure 4-24.

Note: Wait until the volume copy has completed successfully before attempting to use the target volume.

7. Select the copy pair in the Copy Manager, and then select **Copy > Remove Copy Pairs** to remove the copy from the table.

The Read-Only attribute for the target volume is also removed, enabling hosts read and write access.

P

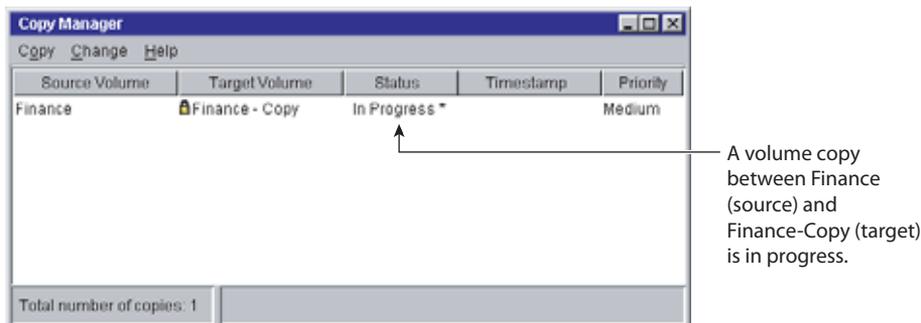


Figure 4-24 Viewing the Volume Copy Status in the Copy Manager

8. Define an additional mapping for the target volume (containing the copied data) by mapping it to the original host, using the LUN recorded in step 4, as shown in Figure 4-25.
 - a. From the **Topology** pane of the **Mappings View**, select the target volume, and then select **Mappings >> Define >> Additional Mapping**.
The **Define Additional Mapping dialog** is displayed.
 - b. Select the same host and LUN used for the source volume, and then choose the target volume from the list of available volumes.
 - c. Select **Add** to define the additional mapping.

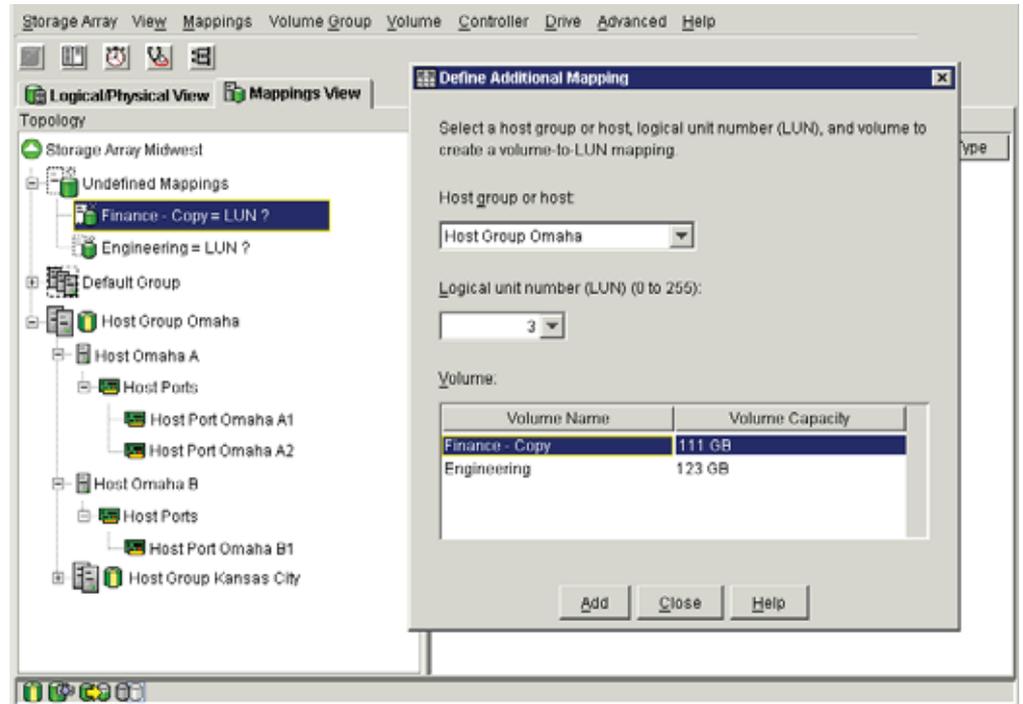


Figure 4-25 Defining an Additional Volume-to-LUN Mapping for the Target Volume

9. Use the `hot_add` utility, or un the following utilities to register the additional volume-to-LUN mapping with the host operating system.

As root, issue the following commands:

- a. # **scsiha -p bus_number**

For example, enter `scsiha -p 2 6 7`

- b. # **ioconfig -f/hw**

This assigns logical controller numbers to all devices in the I/O subsystem.

For more information about the `scsiha` and `ioconfig` utilities, see the `scsiha(1m)` and the `ioconfig(1m)` man pages.

10. Run the `smeedevices` utility to ensure that the volume name and the physical device name (assigned by the operating system) correspond.

11. Remount any file systems on the source volume and target volume.
12. Restart I/O activity to the source volume and target volume.
13. Delete the original source volume, if needed.

It may be desirable to maintain the original source volume for a period of time in case you need to access the data from that volume.

Restoring Snapshot Volume Data to the Base Volume

The Volume Copy premium feature enables you to restore data from a snapshot volume to its original base volume. After a volume copy has been created, it may become apparent, due to possible corruption or data loss, that the data on the snapshot volume is preferred to the data on its associated base volume.

If you need to restore data to the base volume from its associated snapshot volume, the Volume Copy premium feature can be used to copy the data from the snapshot volume to the base volume.

In the following example, on Storage Array Midwest, volume Finance contains business critical data. A point-in-time snapshot copy was created for backup purposes, as shown in Figure 4-26. After the backup was created, data on the original base volume was accidentally deleted or corrupted.

The data from the snapshot volume is now preferred to the data on the base volume, and can be restored to the original base volume using the Volume Copy premium feature.

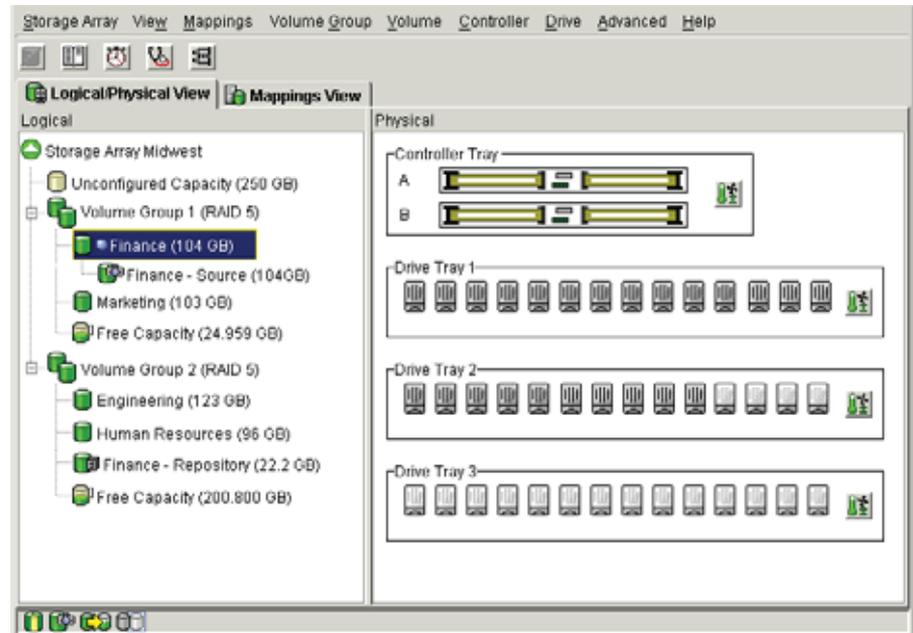


Figure 4-26 Storage Array Midwest Example

To copy the data on an existing snapshot volume back to its associated base volume, complete the following steps.

Note: This procedure assumes that a snapshot volume has already been created.

1. Stop all I/O activity to the base volume.
2. Unmount any file systems on the base volume.
3. If a suitable target volume has not already been created, create a standard volume with a capacity equal to the capacity of the base volume of the snapshot volume, as shown in Figure 4-27.
 - a. Select a volume group's Free Capacity node or the storage array's Unconfigured Capacity node, and then select **Volume >> Create**.

The Create Volume Wizard is displayed.

- b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

Note: To make locating the target volume in the Logical View easier, ensure that you add “target” or “copy” to the volume name. For more information about creating a standard volume, refer to the **Array Management** window online help.

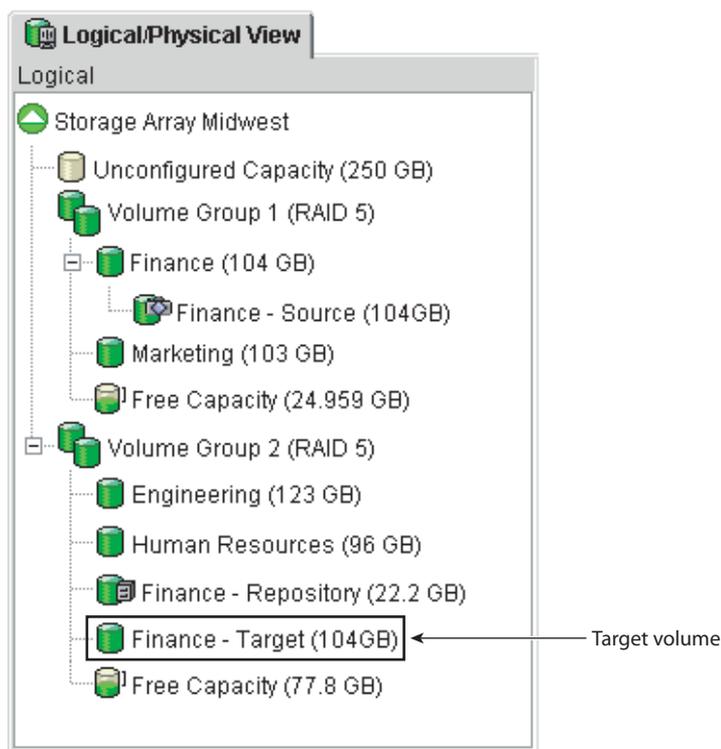


Figure 4-27 Creating Target Volume Finance-Target

- 4. Use the Create Copy Wizard to create a new volume copy between the snapshot volume and the target volume, as shown in Figure 4-28.
 - a. Select the snapshot volume, and then select **Volume >> Copy >> Create**.

The **Create Copy Wizard - Introduction** dialog is displayed, and the source volume is highlighted in the table.

- b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

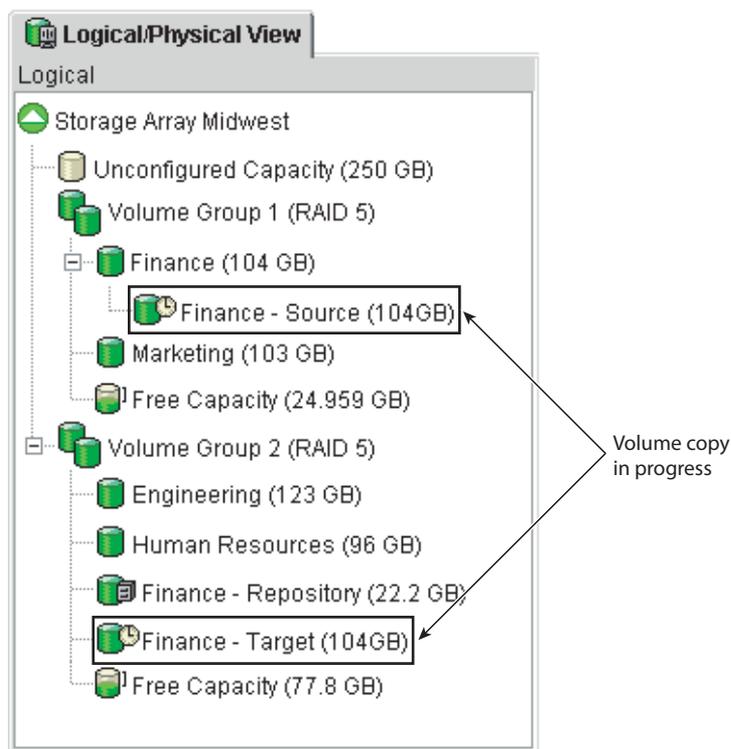


Figure 4-28 Creating a Volume Copy

5. Select **Volume >> Copy >> Copy Manager** to open the Copy Manager to view the status of the volume copy, as shown in Figure 4-29.

Note: Wait until the volume copy has completed successfully before attempting to use the target volume.

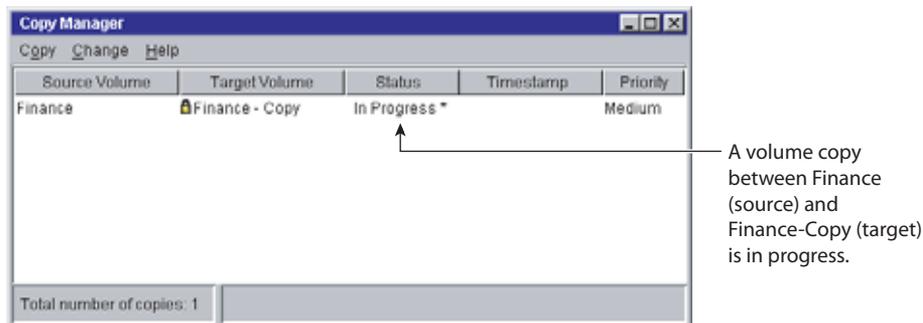


Figure 4-29 Viewing the Volume Copy Status in the Copy Manager

6. Select the copy pair in the Copy Manager and select **Copy >> Remove Copy Pairs** to remove the copy pair from the table.

The Read-Only attribute for the target volume is also removed, enabling hosts to have read and write access to the target volume.

7. Remove the volume-to-LUN mapping for the base volume (Finance). Ensure that you record the LUN number used to provide access to the base volume. You will need this in step 13 when you define a volume-to-LUN mapping for the target volume.
 - a. From the **Mappings View** select the base volume, and then select **Mappings >> Remove**.

A confirmation dialog is displayed.

 - b. Select **Yes**.

The volume-to-LUN mapping is removed, and the base volume is moved back under the Undefined Mappings node in the **Topology View**.

8. Run the following utilities to register the additional volume-to-LUN mapping with the host operating system.

As root, issue the following commands:

- a. # **scsiha -p bus_number**
For example, enter `scsiha -p 2 6 7`
- b. # **ioconfig -f /hw**

This assigns logical controller numbers to all devices in the I/O subsystem.

For more information about the `scsiha` and `ioconfig` utilities, see the `scsiha(1m)` and the `ioconfig(1m)` man pages.

9. Disable the snapshot volume to allow the base volume to become a valid target volume for the volume copy, as shown in Figure 4-30.

Note: If the base volume has multiple associated snapshot volumes, you will need to disable all the associated snapshot volumes before the base volume can be selected as a valid target volume.

- a. Select the snapshot volume in the **Logical View**, and then select **View >> Snapshot >> Disable**.

The **Disable Snapshot Volume** dialog is displayed.

- b. Type **Yes** in the text field and select **OK**.

The snapshot volume is disabled and copy-on-write activity to the snapshot repository volume is stopped until the snapshot volume is recreated.

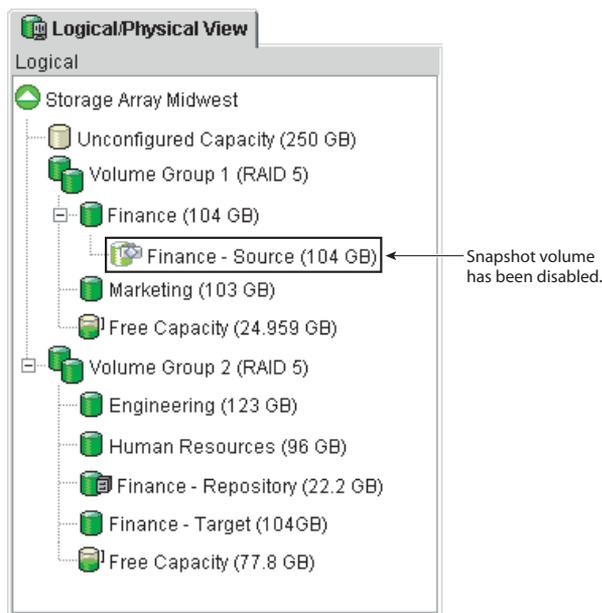


Figure 4-30 Disabling the Snapshot Volume

10. Use the Create Copy Wizard to create a new volume copy between the target volume and the original base volume to restore the data to the base volume (Finance), as shown in Figure 4-31.
 - a. Select the target volume from the previous volume copy, and then select **Volume >> Copy >> Create**.

The **Create Copy Wizard - Introduction** dialog is displayed.

- b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

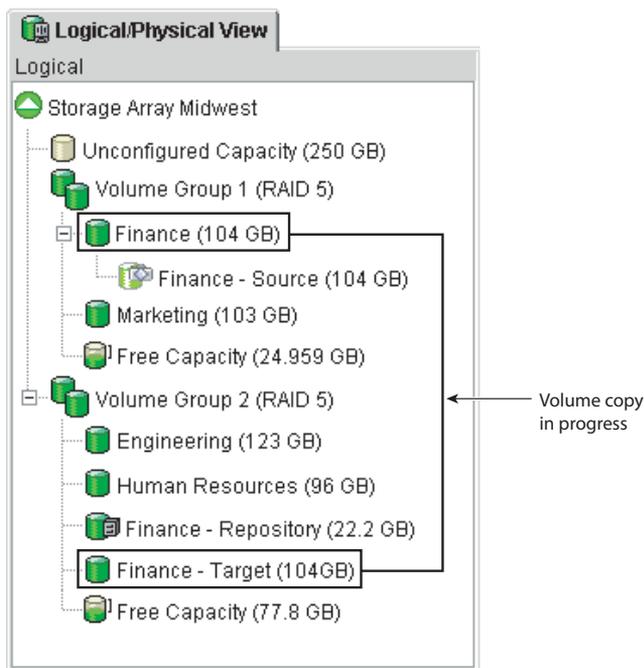


Figure 4-31 Creating a Volume Copy of Finance-Target

11. Select **Volume >> Copy >> Copy Manager** to open the Copy Manager to view the status of the volume copy, as shown in Figure 4-32.

Note: Wait until the volume copy has completed successfully before attempting to use the target volume.

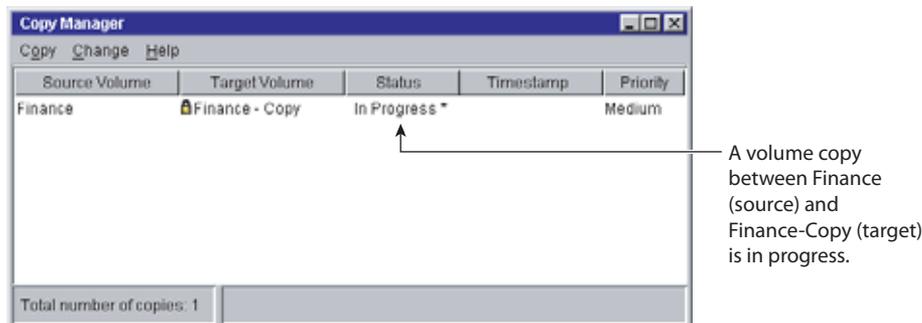


Figure 4-32 Viewing the Volume Copy Status in the Copy Manager

12. Select the copy pair in the Copy Manager and select **Copy >> Remove Copy Pairs** to remove the copy pair from the table.

The Read-Only attribute for the target volume is also removed, enabling hosts to have write access to the target volume.

13. Define a volume-to-LUN mapping for the base volume (containing the copied data) by mapping it to the original host using the LUN recorded in step 7, as shown in Figure 4-33.
 - a. From the **Topology** pane of the **Mappings View**, select the target volume and then select **Mappings >> Define >> Additional Mapping**.
The **Define Additional Mapping** dialog displayed.
 - b. Select the same host and LUN used for the base volume, and then choose the target volume from the list of available volumes.
 - c. Click **Add** to define the additional mapping.

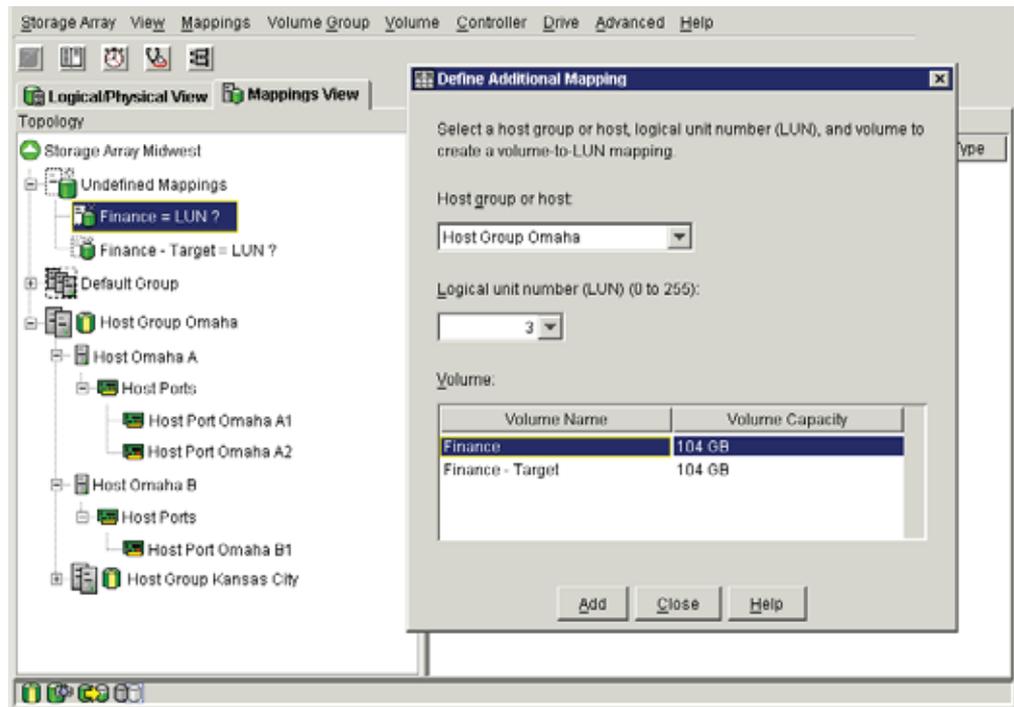


Figure 4-33 Defining an Additional Volume-to-LUN Mapping for the Base Volume

14. Run the following utilities to register the additional volume-to-LUN mapping with the host operating system.

As root, issue the following commands:

- a. # **scsiha -p bus_number**

For example, enter `scsiha -p 2 6 7`

- b. # **ioconfig -f /hw**

This assigns logical controller numbers to all devices in the I/O subsystem.

For more information about the `scsiha` and `ioconfig` utilities, see the `scsiha(1m)` and the `ioconfig(1m)` man pages.

15. Run the `smeedevices` utility to ensure that the volume name and the physical device name (assigned by the operating system) correspond.

For more information about the `smeedevices` utility, refer to the *SGI InfiniteStorage ISSM-EE Administration Guide*.

16. If you need a new point-in-time image of the data on the base volume, recreate the snapshot volume you disabled in step 9.
 - a. Select **Volume >> Snapshot >> Create**.
The **Create Snapshot Volume Wizard - Introduction** dialog is displayed.
 - b. Follow the directions on each Wizard dialog, and select Next when you are ready to move to the next dialog.
17. Remount any file systems on the base volume and target volume.
18. Restart I/O activity to the base volume.

Volume Backup and Restoration

The Volume Copy premium feature allows you to create a backup of a volume by copying data from one volume (source) to another volume (target) in the same storage array, minimizing the time that the source volume is unavailable to host writes.

The target volume can then be used as a backup for the source volume, as a resource for system testing, or to copy data to another device such as a tape drive or other media.

Creating a Backup Copy

The following example details how to back up data on your storage array using the Volume Copy premium feature.

In this example, on Storage Array Midwest, volume Finance contains data that needs to be backed up. The backup will be created on the storage array, and eventually transferred to a tape drive for off-site storage.

The source volume's actual capacity is 104 GB, as shown in Figure 4-34. Using the Volume Copy premium feature, a complete copy of Finance will be created and stored on the storage array as a backup of the base volume.

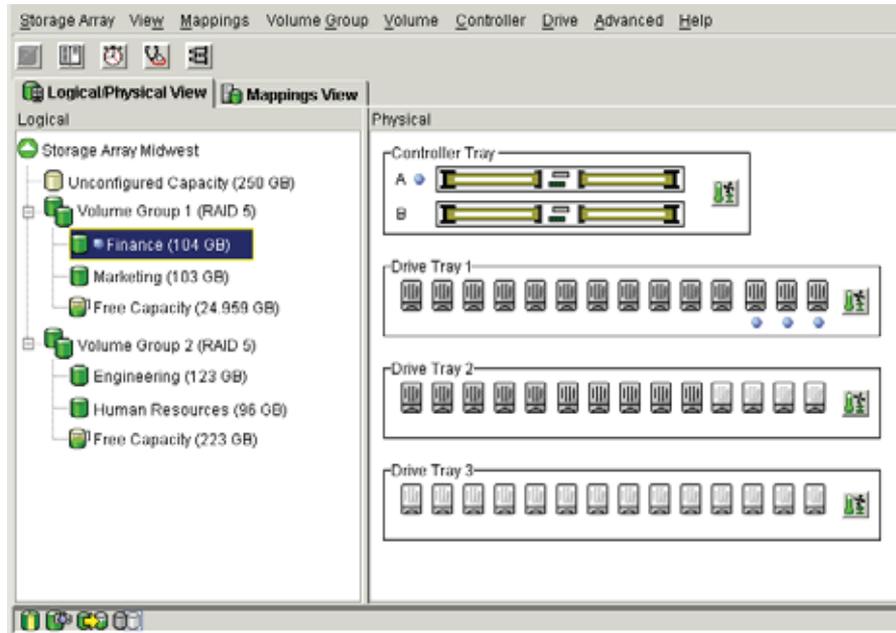


Figure 4-34 Storage Array Midwest Example

To copy data from the base volume to the target volume, which can then be used to back up data to a drive tape or some other media, complete the following steps.

1. Stop all I/O activity to the base volume (Finance).
2. Unmount any filesystems on the base volume.
3. Select the base volume (Finance) in the **Logical View** of the **Array Management** window.
4. Use the Create Snapshot Volume Wizard to create a snapshot volume, as shown in Figure 4-35.
 - a. Select **Volume >> Snapshot >> Create**.
The **Create Snapshot Volume Wizard - Introduction** dialog is displayed.
 - b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

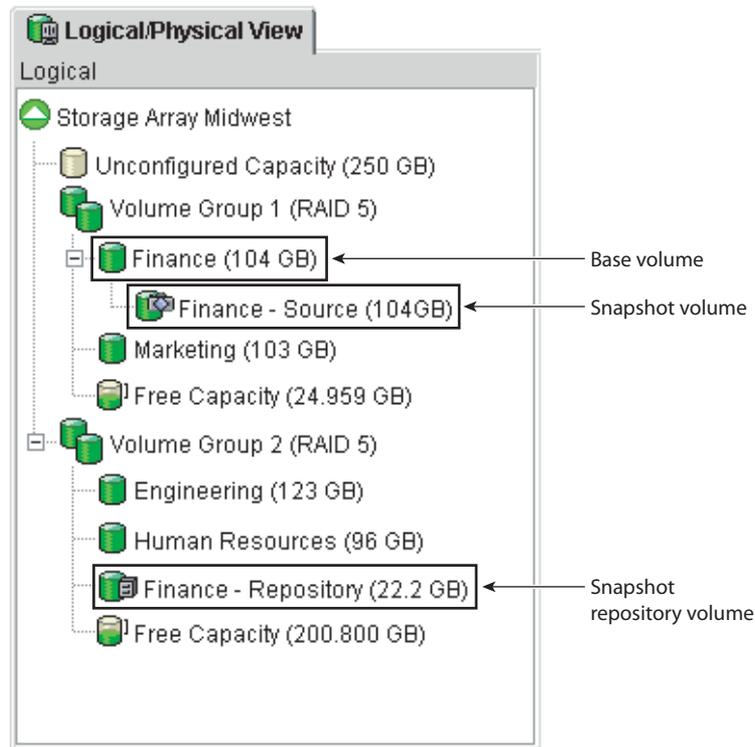


Figure 4-35 Creating a Snapshot Volume of Finance

5. Remount any file systems on the base volume.
6. Restart I/O activity to the base volume.

Finance is now available for host read and write access. The snapshot volume created in step 4 can now be selected as the source volume for the volume copy.

7. If a target volume has not already been created, create a standard volume with a capacity equal to or greater than the snapshot volume's capacity, as shown in Figure 4-36.
 - a. Select a volume group's Free Capacity node or the storage array's Unconfigured Capacity node, and then select **Volume >> Create**.

The Create Volume Wizard is displayed.

- b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

Note: To make locating the target volume in the Logical View easier, ensure that you add “target” or “copy” to the target volume name. For more information about creating a standard volume, refer to the **Array Management** window online help.

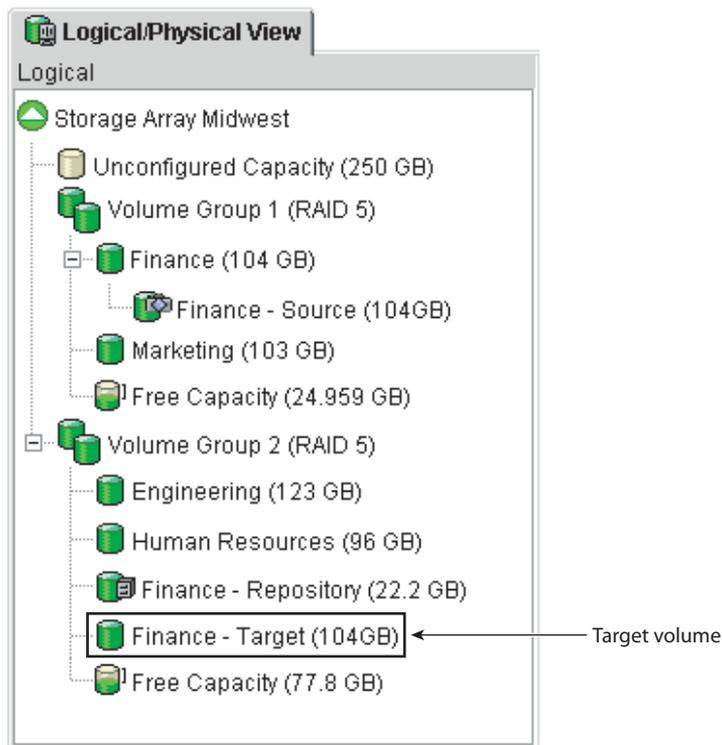


Figure 4-36 Creating the Target Volume Finance-Target

- 8. Use the Create Copy Wizard to create a new volume copy, using the snapshot volume as the source volume and the volume you created in step 7 as the target volume, as shown in Figure 4-37.
 - a. Select the snapshot volume, then select **Volume >> Copy >> Create**.

The **Create Copy Wizard - Introduction** dialog is displayed, with the snapshot volume highlighted in the table.

- b. Follow the directions on each Wizard dialog, and select **Next** when you are ready to move to the next dialog.

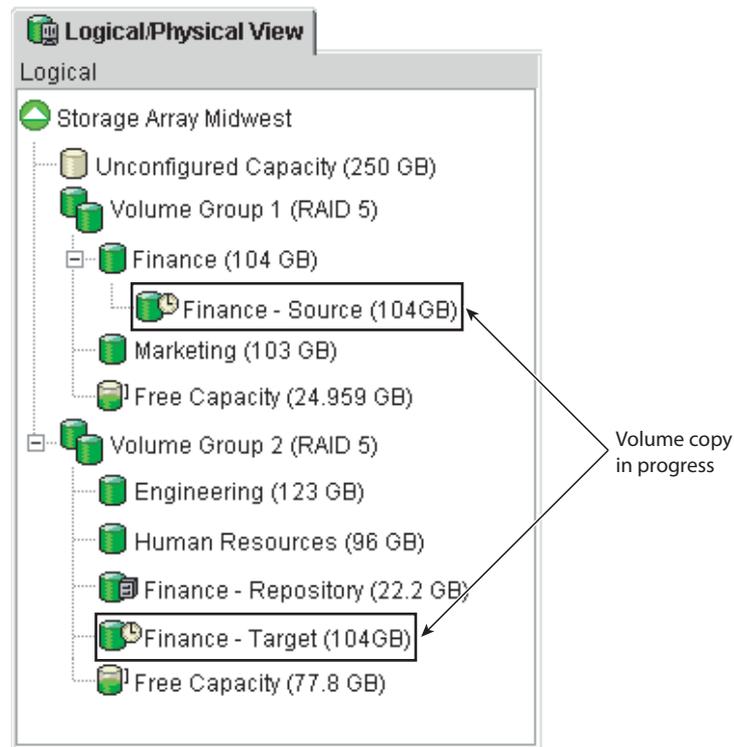


Figure 4-37 Creating a Volume Copy the Snapshot Volume Finance-Copy

9. Select **Volume >> Copy >> Copy Manager** to open the Copy Manager to view the status of the volume copy, as shown in Figure 4-38.

Note: Wait until the volume copy has completed successfully before attempting to use the target volume.

10. Select the copy pair in the Copy Manager, and select **Copy >> Remove Copy Pairs** to remove the copy pair from the table.

The Read-Only attribute for the target volume is also removed, enabling hosts to have write access to the target volume.

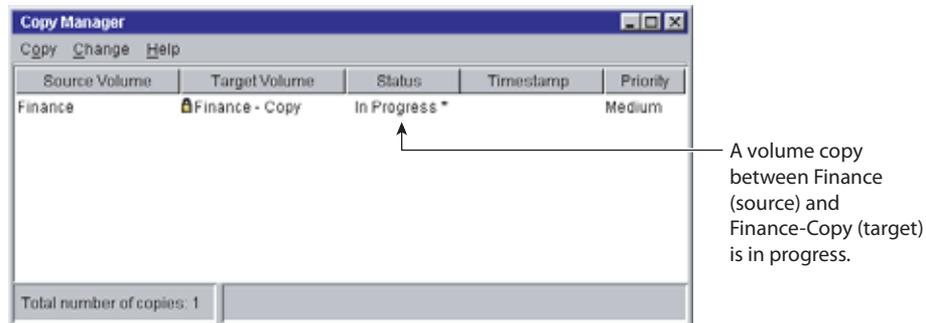


Figure 4-38 Viewing the Status of the Volume Copy in the Copy Manager

11. Disable the snapshot volume to stop any further copy-on-write activity to its associated snapshot repository volume, as shown in Figure 4-39.
 - a. Select the snapshot volume in the Logical View, and then select **Volume >> Snapshot >> Disable**.
The **Disable Snapshot Volume** dialog is displayed.
 - b. Type **Yes** in the text field, and then select **OK**.
The snapshot volume is disabled, and copy-on-write activity to the snapshot repository volume is stopped until the snapshot volume is recreated.

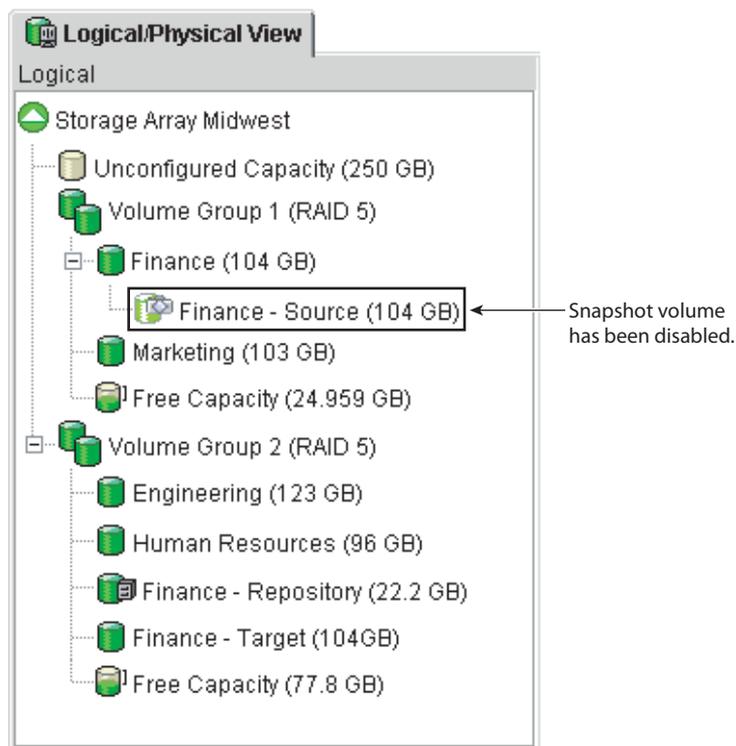


Figure 4-39 Disabling Snapshot Volume Finance-Source

12. The target volume can now be used to back up data to a tape drive (to be stored off-site), for application or system testing, or be retained on the storage array for quick recovery of the data on the base volume (Finance).

Recovering from a Backup Copy

Data can be recovered instantly from the backup volume you created in the previous procedure by simply unmapping the base volume (Finance) from its host, and then mapping the backup volume (Finance - Target) to the same host.

To recover data from Finance - Target, complete the following steps.

1. Stop all I/O activity to the base volume (Finance).
2. Unmount any file systems on volume Finance.
3. Remove the volume-to-LUN mapping for the original base volume (Finance). Ensure that you record the LUN number used to provide access to the base volume. You will need this in step 4 when you define a volume-to-LUN mapping for the target volume.
 - a. From the **Mappings View** select the base volume, and then select **Mappings >> Remove**.
A confirmation dialog is displayed.
 - b. Select **Yes**.
The volume-to-LUN mapping is removed, and the base volume is moved back under the Undefined Mappings node in the **Topology View**.
4. Run the following utilities to register the additional volume-to-LUN mapping with the host operating system.
As root, issue the following commands:
 - a. # **scsiha -p bus_number**
For example, enter `scsiha -p 2 6 7`
 - b. # **ioconfig -f /hw**
This assigns logical controller numbers to all devices in the I/O subsystem.For more information about the `scsiha` and `ioconfig` utilities, see the `scsiha(1m)` and the `ioconfig(1m)` man pages.
5. Define a volume-to-LUN mapping for the target volume (containing the copied data), using the same LUN as the original base volume, recorded in step 3.
 - a. From the **Topology** pane of the **Mappings View**, select the target volume and then select **Mappings >> Define >> Additional Mapping**.

Tip: The **Define Additional Mapping** dialog is displayed.

- b. Select the same host and LUN used for the base volume, and then choose the target volume from the list of available volumes.
 - c. Click **Add** to define the additional mapping.
6. Run the following utilities to register the additional volume-to-LUN mapping with the host operating system.

As root, issue the following commands:

- a. # **scsiha -p bus_number**

For example, enter `scsiha -p 2 6 7`

- b. # **ioconfig -f /hw**

This assigns logical controller numbers to all devices in the I/O subsystem.

For more information about the `scsiha` and `ioconfig` utilities, see the `scsiha(1m)` and the `ioconfig(1m)` man pages.

7. Run the `smeedevices` utility to ensure that the volume name and the physical device name (assigned by the operating system) correspond.

For more information about the `smeedevices` utility, refer to the *SGI InfiniteStorage ISSM-EE Administration Guide*.

8. Remount any file systems on the target volume.
9. Restart I/O activity to the base volume and target volume.

Troubleshooting

This section provides information about troubleshooting some of the more common problems encountered when using the Volume Copy premium feature.

Troubleshooting topics included are general Volume Copy problems and premium feature troubleshooting.

Locate the section that is directly related to your problem, or consult the section describing general troubleshooting techniques. If your problem is not discussed in any of the sections in this chapter, contact technical support.

General Troubleshooting

This section describes some of the typical problems you may encounter when using the Volume Copy premium feature. The information in this section is presented in sequence, from the time the premium feature is enabled to problems that you may encounter using the Volume Copy premium feature.

Information about the probable cause and recommended resolution for each problem are provided. If your problem is not addressed in this chapter, contact technical support.

The Volume Copy Menu options in the Array Management window are greyed out and unavailable

Cause: The Volume Copy premium feature is not enabled.

Resolution: Enable the Volume Copy premium feature. You must use a Feature Key File to enable premium features. If one has not been supplied, contact your storage supplier.

For step-by-step instructions about enabling the Volume Copy premium feature, refer to “Enabling the Volume Copy Premium Feature” on page 195 or the **Array Management** window online help.

The Volume Copy icon in the premium status area is greyed out with a red line

Cause: The Volume Copy premium feature is not enabled.

Resolution: Enable the Volume Copy premium feature. You must use a Feature Key File to enable premium features. If one has not been supplied, contact your storage supplier.

For step-by-step instructions about enabling the Volume Copy premium feature, refer to “Enabling the Volume Copy Premium Feature” on page 195 or the **Array Management** window online help.

The Create Copy option is not available after selecting a source volume in the Logical View

Cause: The volume you have chosen to be the source volume for the volume copy is not a valid source volume.

The following volumes are not valid source volumes and will not appear in the list:

- Remote Volume Mirror secondary volume
- Snapshot repository volume
- Remote Volume Mirror repository volume
- Failed volume
- Missing volume
- A volume currently in a modification operation
- A volume that is reserved by the storage management software
- A volume that is a source volume or target volume in another volume copy with a status of In Progress, Failed, or Pending

Resolution: Select a valid source volume, which can be a standard volume, snapshot volume, the base volume of a snapshot volume, or a primary volume participating in a Remote Volume Mirror.

The source volume/target volume is displayed as being a Missing Volume

Cause: A physical drive on which the Volume Copy depends (source volume or target volume) has been removed, and the volume is now displayed under the Missing Volumes node in the **Logical View** of the **Array Management** window.

Resolution: If missing volumes are detected because drives have accidentally been removed, or are detected as missing because of a power loss to the drive modules, recovery of these volumes is possible by completing the following steps:

1. Re-insert the drives into the drive module.
2. Ensure that the drive module's power supplies are properly connected to an operating power source and have Optimal status.

For more information about troubleshooting missing volumes, refer to the Recovery Guru or the **Array Management** window online help or contact technical support. If the missing volume is no longer required, it may be deleted.

Caution: Deleting a missing volume is a permanent action. Any associated volumes or volume-to-LUN mappings will also be deleted.

The controller owner for the source volume/target volume changes after the volume copy has completed

- Cause:** When the volume copy started, the preferred controller did not have ownership of both the source volume and target volume.
- Resolution:** During a volume copy, the same controller must own both the source volume and the target volume. If both volumes do not have the same preferred controller when the volume copy starts, the ownership of the target volume is automatically transferred to the preferred controller of the source volume.
- When the volume copy is completed or is stopped, ownership of the target volume is restored to its preferred controller. If ownership of the source volume is changed during the volume copy, ownership of the target volume is also changed.

A volume copy that includes a source volume participating in a Remote Volume Mirror has a status of In Progress, Pending, or Completed, but has failed

- Cause:** In response to a catastrophic failure at the primary storage array, a role reversal has occurred. The secondary volume in the mirrored pair was promoted to the primary volume role, while the primary volume was demoted to the secondary volume role.
- Resolution:** Once the preferred controller of the primary volume can be contacted, a role reversal will automatically take place and you can then restart the volume copy.
- For more information about role reversals, refer to "Role Reversals" on page 193.

Error message indicates that the Volume Copy premium feature is out of compliance

- Cause:** The Volume Copy Premium Feature is out of compliance.
- This normally occurs if a volume group with existing Volume Copies has been moved into a storage array where the feature is not authorized, or if you have disabled the feature on a storage array where the feature is authorized, and you have existing volume copies present.
- Resolution:** If the Volume Copy feature is not authorized for this storage array and you want to authorize it, or the Volume Copy feature is authorized for this storage array but you have disabled it, complete the following steps to clear the out of compliance error:
1. Obtain a new Feature Key File for this feature, or retrieve the existing Feature Enable Identifier by selecting **Storage Array >> Premium Features >> List** in the **Array Management** window.
 2. Enable the Volume Copy premium feature. For step-by-step instructions, see “Enabling the Volume Copy Premium Feature” on page 195 or the **Array Management** window online help.
- If the Volume Copy feature is not authorized for this storage array and you do not want to authorize it, disable the Volume Copy premium feature by selecting the **Storage Array >> Premium Features >> Disable** option in the **Array Management** window.
- For more information about disabling premium features, refer to the **Array Management** window online help.

Error message indicates that the selected source volume has a reservation placed on it

- Cause:** Volumes configured with a legacy or persistent reservation cannot be selected as a source volume or a target volume for a volume copy.
- Resolution:** Clear the reservation for the selected volume, and then recreate the volume copy.
1. Select **Advanced >> Persistent Reservations** to open the **Persistent Reservations** dialog.
 2. Ensure that the **View Associated Registrations** check box is unselected in the upper left corner of the **Persistent Reservations** dialog.

3. Click one or more desired volumes. To select all volumes, click **Select All**. The selected volumes or all volumes are highlighted, depending on your selection.
4. Select **Clear**.
The **Clear Registrations/Reservations** dialog is displayed.
5. Type **Yes** in the confirmation field.
The **OK** button becomes enabled.
6. Select **OK**, or select **Cancel** to return to the main **Persistent Reservations** dialog without clearing any reservations.

The reservations and registrations associated with the volumes that you have highlighted in the upper pane are cleared.

Volume Copy Critical Events

Critical Major Event Log (MEL) events are generated when an error occurs that can affect data availability or that results in a degraded mode of operation. The storage management software responds by alerting the appropriate administrators via email, SNMP trap, or other configured mechanisms.

Because administrative action is required to correct the problem, the storage array will generally enter a Needs Attention status. This causes the appropriate indications to be presented through the storage management software and an associated Recovery Guru procedure to be displayed when requested.

The following table provides a description of critical MEL events, with a brief explanation of what caused the event, and the recommended resolution. If your critical MEL event is not included in this section, contact technical support.

Event 6600 - Volume Copy Operation Failed

Cause: This event is logged when a Volume Copy with a status of In Progress or Pending has failed. A Needs Attention status is also reported for the storage array.

This failure can be caused by a read error on the source volume, a write error on the target volume, or because of a failure that occurred on the storage array that affects the source volume or target volume.

Resolution: For specific recovery procedures, refer to Recovery Guru.

Glossary of Volume Copy Terms

This section provides definitions of terms that you may encounter when using the Volume Copy feature.

Array Management window (AMW)

The interface from which you can manage a storage array. Use the **Array Management** window to manage and configure physical and logical components, recover from errors, and monitor performance.

Base Volume A volume from which you create a snapshot volume. A base volume must be a standard volume in the storage array.

Copy Complete

This status is displayed in the Copy Manager signifying that the data on the source volume has been successfully copied to the target volume. This status is accompanied by a timestamp attribute.

Copy Failed This status is displayed in the Copy Manager when an error occurred during the volume copy. A status of Failed can occur because of a read error from the source volume, a write error to the target volume, or because of a failure on the storage array that affects the source volume or target volume.

A critical event is logged in the Major Event Log (MEL) and a Needs-Attention icon is displayed.

Copy-In-Progress

This status is displayed in the Copy Manager while data on the source volume is being read and then written to the target volume. While a volume copy has this status, the host has read-only access to the source volume, and read and write requests to the target volume will not take place until the volume copy has completed.

Copy Manager

The Copy Manager is a dialog in the Volume Copy premium feature used to monitor the progress of a volume copy, and to manage all volume copies on a single storage array after they have been created.

Copy-on-Write

A copy-on-write operation occurs the first time any data block is modified on the base volume of a snapshot volume. The point-in-time data you wish to preserve from the snapshot volume is stored on the base volume until the storage array controller receives a write request

for the base volume. Before the write occurs, the data blocks to be modified are copied over to the snapshot repository volume for safekeeping.

Copy Pair A copy pair refers to a source volume and a target volume in a volume copy that are located on the same storage array.

Copy Pending A volume copy status displayed in the Copy Manager when a volume copy has been created, but system resources do not allow it to start. While in this status, the host has read-only access to the source volume, and read and write requests to the target volume will be rejected until the volume copy has completed.

Copy Priority This setting defines how much of the system resources are used to complete a volume copy versus fulfill I/O requests. There are five relative settings ranging from Lowest to Highest. The Highest priority rate supports the volume copy, but I/O activity may be affected. The Lowest priority rate supports I/O activity, but the volume copy will take longer.

You can set the copy priority for each volume copy in the Create Copy Wizard or by using the Copy Manager.

Copy Stopped This status is displayed in the Copy Manager immediately after a volume copy is created in the Wizard, or when the volume copy is stopped in the Copy Manager. This status does not occur as a result of a failure.

While a volume copy has this status, no data is copied between the source volume and target volume.

Enterprise Management window (EMW)

A graphical user interface (GUI) from which a user can manage storage arrays, monitor the health of the storage arrays, configure e-mail and Simple Network Management Protocol (SNMP) alert notifications, launch an appropriate **Array Management** window for a particular storage array, or edit and execute a script to perform batch management tasks.

Feature Key File

The storage array controller uses this file to enable an authorized premium feature. The file contains the Feature Enable Identifier of the storage array for which the premium feature is authorized, and contains generated data about the premium feature.

Missing Volume

A missing volume is a placeholder node displayed in the **Logical View**, indicating that the storage array has detected that drives associated with a volume are no longer accessible. Typically, this is the result of removing drives associated with a volume group or the loss of power to one or more drive trays.

Persistent Reservation

A reservation on a volume that prevents other hosts from accessing the volume. Unlike other types of reservations, a persistent reservation reserves access across multiple host ports, provides various levels of access control, offers the ability to query the storage array about registered ports and reservations, and optionally, provides for persistence of reservations in the event of a storage array power loss.

Preferred Controller Ownership

The preferred controller ownership of a volume or volume group is the controller of an active-active pair that is designated to own these volumes. The preferred controller owner is the controller that currently owns the volume or volume group.

If the preferred controller is being replaced or is undergoing a firmware download, ownership of the volumes is automatically shifted to the other controller, and that controller becomes the current owner of the volumes. This is considered a routine ownership change and is reported with an informational entry in the event log.

There can also be a forced failover from the preferred controller to the alternate controller because of I/O path errors. This is reported with a critical entry in the event log, and will be reported by the Enterprise Management software to e-mail and SNMP alert destinations.

Premium Feature

A storage array feature that may not be available in the standard configuration of the storage management software.

Use the **Storage Array >> Premium Features >> Enable** pull-down menu option to enable a premium feature.

Primary Volume

The primary volume in a Remote Volume Mirror is the volume that accepts host I/O and stores application data. The data on a primary volume is replicated to the secondary volume.

Secondary Volume

A secondary volume is used to store data replicated from a primary volume. The controller owner of a secondary volume receives remote writes for the volume from the controller owner of the primary volume. The controller owner of the secondary volume does not accept host read or write requests.

smeedevices

`smeedevices` is a host-based utility used to associate the mapping between a physical device name and a volume name.

Once volumes have been created and volume-to-LUN mappings have been defined, this utility is run to ensure that the volume name and the physical device name (assigned by the operating system) correspond.

Snapshot Repository Volume

A snapshot repository volume is a special volume in the storage array that is created as a resource for a snapshot volume. A snapshot repository volume contains snapshot data and copy-on-write data for a particular snapshot volume.

A snapshot repository volume is not host-addressable, but you can increase its capacity, rename it, specify the action (either failing it or failing writes to the base volume) when it becomes full, or view its capacity or media scan properties.

Snapshot Volume

A snapshot volume is a logical point-in-time image of another volume, called a base volume, in the storage array. A snapshot is the logical equivalent of a complete physical copy, but you create it much more quickly than a physical copy and it requires less disk space.

Because a snapshot is a host-addressable volume, you can perform backups using the snapshot while the base volume remains online and user-accessible. In addition, you can write to the snapshot volume to perform application testing or scenario development and analysis.

Source Volume

A standard volume in a volume copy that accepts host I/O and stores application data. When the volume copy is started, data from the source volume is copied in its entirety to the target volume.

Standard Volume

A logical structure you create on a storage array for data storage.

Storage Array A storage entity managed by storage management software. A storage array consists of a collection of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as volume groups and volumes). A storage array can span multiple physical enclosures.

For example, several drive modules connected to a command module constitutes one storage array. A unit that contains drives and a controller in a single enclosure is also considered one storage array. A rackmount subsystem may contain more than one storage array.

Also known as a RAID Module.

Target Volume A standard volume in a volume copy that maintains a copy of the data from the source volume. Read and write requests to the target volume will be rejected while the volume copy has a status of In Progress, Pending, or Failed.

When a volume copy is started, the data from the source volume is written to the target volume.

A volume copy will overwrite data on the target volume. Ensure that you no longer need the data on the volume, or have backed it up, before starting the volume copy.

Usable Capacity

The minimum of the primary volume and secondary volume actual capacities.

In a given Array Management window, all primary or secondary volumes present in the same storage array display both actual and usable capacity in the Logical View.

Volume

A volume is a logical structure that you create on a storage array for data storage. A volume is defined over a set of drives called a volume group, and has a defined RAID level and capacity. The drive boundaries of the volume group are hidden to the host computer.

Volume Copy

The Volume Copy premium feature is used to copy data from one volume (the source volume) to another volume (the target volume) within a single storage array.

Volume Reservation

A feature of cluster software that allows one or more host ports to reserve a volume and prevent other host ports from accessing the same volume.

Index

A

access volume, 7, 8

Add Device option, 13

agent software

 see host-agent software

alert

 destinations, 55

 notifications, 11

Array Management window

 described, 17

 failure notification in, 64

 functions, 17

 Help, vii

 introduction to, 9

 monitoring Storage Arrays with, 55

associated volumes, locating, 206

Automatic Discovery option, 13

B

backup copy, creating a, 233

battery age clock, 34

BOOTP or BOOTP-Compliant DHCP Server, 2

C

cache

 block size, 70

 flush

 start, 70

 stop, 70

 settings, 75

cache flush

 described, 34

 performance impacts, 34

 settings, 34

 start percentage, 34

 stop percentage, 34

cache hit percentage

 defined, 70

 impact from cache flush settings, 34

cache read-ahead

 choosing a multiplier, 72

 defined, 72

Campus Configuration

 cabling, 122-125

 defined, 91

 redundancy characteristics, 104, 111

Change Copy Priority option, 214

Components button, 20

configuration

 automatic, 75

 resetting a, 75

configuration file

 described, 11

 recovering a corrupted, 76

Contacting Device status, 58

controller, 44

 defined, 3

 listing components associated with, 77

- ownership, 75
- placing in active or passive mode, 77
- placing online or offline, 77
- viewing properties, 78
- controller host side I/O port
 - and read/write requests, 91
 - dedicated for mirroring, 91, 137, 142, 156
 - requirements, 91
 - zoning ports reserved for mirroring, 114
- controller module
 - see controller tray
- controller owner
 - and full synchronization, 80
 - and remote writes, 80
 - I/O path errors, 93
 - managing data replication, 83
 - of secondary volume, 80
 - prioritization of full synchronization, 89
 - read requests, 84
- controller ownership, 186
- controller resets, 81
- controller tray
 - batteries, 35
 - defined, 3
- Copy Manager, 208
 - Change Copy Priority option, 214
 - Disable Read-Only option, 216
 - Enable Read-Only option, 216
 - Re-Copy option, 210
 - Remove Copy Pairs option, 213
 - Stop Copy option, 212
- copy pair, description of, 184, 208
- copy priority, 187
 - affect on storage array performance, 187
 - changing the, 214
 - description of, 187
 - relative settings for, 187
 - setting the, 187
- Copy Started dialog, 204

- copyback, 35
- copying data, 183
 - for improved access, 218
 - to restore snapshot volume data, 224
- Create Copy Wizard
 - description of, 185
 - dialogs in the, 185
 - using the, 201
- Create Remote Volume Mirror Wizard, 143
- critical alerts, 15
- critical data, backing up, 158
- critical events, 246
- Critical Major Event Log (MEL), 179, 246

D

- data protection strategies, 29
 - Hardware Redundancy, 34
- data replication
 - and link interruptions, 83
 - description of, 83
 - process illustration, 84
- Default Group, defined as a topological element, 44
- default LUN, 44
- default volume-to-LUN mapping
 - and volumes, 51
- Defined Mappings view, 44
- device
 - automatic discovery of, 73
 - correcting a partially managed, 76
 - manually adding, 73
 - removing, 73
- Device Table, 57
- Device Tree, 57
- directly managed storage array, 7
- disable target volume read-only attribute, 216
- Disabled/Activated status

- changing the status, 157
 - defined, 135
 - Disabled/Deactivated status
 - defined, 135
 - icon, 133
 - limitations of, 156
 - disaster recovery
 - backing up critical data, 158
 - preparation steps, 158
 - drive
 - defined, 3
 - failing, 76
 - initializing, 77
 - listing components associated with, 77
 - locating, 77
 - reconstructing, 77
 - reviving, 77
 - viewing properties, 78
 - drive group
 - see volume group
 - drive module
 - see drive tray
 - drive tray
 - defined, 3
 - drives, viewing information for, 200
 - dual-primary condition, 165, 194
 - dual-secondary condition, 166
 - Dynamic Volume Expansion (DVE), 99
- E**
- EL, 246
 - email destinations, 76
 - emwdata.bin file, 11
 - enable target volume read-only attribute, 217
 - Enabled/Activated status
 - changing the status, 157
 - defined, 136
 - Enabled/Deactivated status defined, 136
 - enabling the premium feature, 195
 - Enterprise Management window
 - figure, 12
 - functions, 11
 - introduced, 9
 - monitoring storage arrays with, 55
 - status icons displayed in, 57
 - Enterprise Management window Help, vii
 - entire site failover
 - description of, 161
 - Event Log, 76
 - Event Monitor
 - setting up alert destinations, 60
 - synchronizing with Enterprise Management window, 61
 - expected usage, of a volume, 42
- F**
- failure notification
 - in the Array Management window, 64
 - types of, 63
 - Feature Key File, 136, 195
 - Fibre Channel fabric configurations
 - defined, 91
 - distance limit, 94
 - Fibre Channel switches. See switches
 - finding associated volumes, 206
 - firmware
 - downloading, 74
 - firmware version
 - requirements, 101
 - Fixing status, 57
 - free capacity, 4
 - adding to a volume group, 42, 74
 - consolidating on a volume group, 74

- creating volumes from, 74
- creating volumes on, 42
- defined, 4
- increasing storage array, 74

Free Capacity node, 4

full synchronization

- after connectivity is restored, 87
- after mirror relationship is created, 83
- defined, 80
- performance considerations, 89
- time to complete, 89

G

general troubleshooting, 242

Go To option, 206

H

hardware components, 2

Help

- task reference, 73
- using, vii

heterogeneous hosts, overview and example, 52

Highest Availability Campus Configuration

- cabling, 115-119
- defined, 91
- redundancy characteristics, 102, 111

host

- defined, 2

host bus adapter utility, 45

host cluster software, 163

host group, defined as a topological element, 44

host port, defined as a topological element, 45

host type, setting, 53

host, defined as a topological element, 44

host-agent managed Storage Array

- described, 7
- discovering newly attached, 73

host-agent software

- described, 8
- where installed, 8

hot spare drive

- assigning, 74
- configuring, 35
- defined, 35
- unassigning, 74

hot_add utility, 223

I

indicator lights, 77

Intra-Site Configuration

- cabling, 131
- defined, 92
- redundancy characteristics, 105, 111, 125

I/O access pattern, 71

I/O data connection

- see I/O data path

I/O data path

- described, 5
- protection, 37

I/O path error, 93

I/O request rate

- defined, 70
- impact from cache flush settings, 34
- modification priority effect on, 73

L

link interruptions, 83

locating associated volumes, 206

logical unit number

- see LUN

- Logical view, 57
 - status icon changes, 67, 68
 - LUN
 - address space, 51
 - defined
 - using in a volume-to-LUN mapping, 51
 - see volume
 - LUN (logical unit number), defined, 45
- M**
- Major Event Log (MEL), 179
 - management domain
 - adding host-agent managed storage arrays to, 7
 - configuration file, 11
 - defined, 5
 - global e-mail alert settings, 76
 - populating, 13
 - use of Automatic Discovery, 13
 - management station
 - see storage management station, 1
 - managing volume copies, 208
 - Mappings View, overview, 43
 - media scan
 - changing settings, 75
 - defined, 36
 - medical imaging applications, 32
 - MEL. See Major Event Log
 - MEL (Major Event Log), 246
 - mirror relationship. See Remote Volume Mirroring
 - mirror relationships
 - maintaining, 145
 - maximum per storage array, 80
 - prerequisites for creating, 83
 - removing, 154, 172
 - resuming, 153
 - suspending, 147
 - suspending in asynchronous mode, 149
 - viewing details, 145
 - mirror repository volume
 - defined, 81
 - deleting, 156
 - information logged to, 83
 - RAID level of, 82
 - size of, 82
 - mirror repository volumes
 - upgrading, 139
 - mirrored volume pair, 191
 - missing volume, primary or secondary, 171
 - modification operations, 72, 195
 - modification priority
 - changing, 75
 - described, 72
 - effect on I/O request rate, 73
 - setting, 73
 - monitoring progress, 208
 - multimedia applications, 32
 - multi-user environments, 33
- N**
- Needs Attention status, 57
 - network management connection
 - combinations of, 6
 - defined, 6
 - host-agent, 8
 - types of, 6
 - Network Management Station
 - defined, 2
 - receipt of failure notifications, 64
 - NVSRAM, downloading, 74
- O**
- online Help

- see Help
- Optimal status, 57, 180
- orphan mirror, 172
- Overall Health Status, 57

P

- password protection, 36, 77
- performance considerations, Remote Volume Mirrors, 89
- Performance Monitor, 70, 75
- persistent reservations, 187
- Physical view
 - changing tray order in, 74
 - described, 20
 - status icon changes, 67, 68
- power outage, 35
- preferred controller owner, 186
- premium features
 - Remote Volume Mirroring, 191
 - Snapshot Volume, 191
 - Storage Partitioning, 189
- Preview dialog, 204
- primary volume, 191
 - and full synchronization, 80
 - base volume for a snapshot, 95
 - controller ownership changes, 93
 - creating, 142
 - defined, 80
 - ownership of, 92
 - role reversals, 193
- properties, Volume Copy, 207

R

- RAID 0
 - described, 32

- drive failure consequences, 32
- RAID 1
 - described, 32
 - drive failure consequences, 32
 - performance, 71
 - redundancy, 30
- RAID 10
 - redundancy, 30
- RAID 3
 - described, 32
 - drive failure consequences, 32
 - redundancy, 30
- RAID 5
 - described, 33
 - drive failure consequences, 33
 - performance, 71
 - redundancy, 30
- RAID 6
 - described, 33
 - drive failure consequences, 33
 - redundancy, 30
- RAID level
 - and drive number constraints, 138
 - changing, 71, 75
 - described, 30
 - for secondary volume, 142, 143
 - of volume group for mirror repository volumes, 138
 - use on a volume group, 30
 - using an appropriate, 30
- RAID Module
 - see Storage Array
- read percentage, 71
- reconstruction, 35
- Re-Copy option, 210
- Recovery Guru
 - described, 65
 - identifying when to use, 76
 - Recovery Procedure, 66

Summary Area, 66

redundancy

- checking, 76
- data, 30
- of Storage Array hardware, 29

remaining capacity

- see free capacity

Remote Volume Mirror. See mirror relationships

Remote Volume Mirroring

- description of, 191
- full synchronization, 191
- primary volume, 191
- role reversal, 193
- secondary volume, 191

remote volume mirroring

- deactivating, 156

Remote Volume Mirroring feature

- and other premium features, 94
- firmware requirements, 101
- hardware requirements, 101
- introduction, 80
- out of compliance, 173
- software installation, 131
- software requirements, 101
- storage management software requirements, 101
- troubleshooting, 167
- verifying configuration before beginning, 132-133

remote write, 83

Remove Copy Pairs option, 213

removing copy pairs, 213

restoring snapshot volume data, 224

restrictions, 194

role reversal

- description of, 193
- performing a, 193
- with snapcopy volumes, 98
- with snapshot volumes, 97
- with Storage Partitioning, 94

S

Script Editor, 11, 74

secondary storage array, 191

secondary volume, 191

- and host reads and writes, 80
- creating, 142
- defined, 80
- determining controller ownership, 92
- failed, 88
- minimum size, 142
- offline, 88
- Optimal status, 88
- preparing for disaster recovery, 159
- RAID level, 142, 143
- restrictions with snapshot premium feature, 96

segment, 71

segment size

- changing, 71, 75
- selecting an appropriate, 71

site preparation, 108

SMrepassist, 205

snapcopy volume

- failure due to role reversal, 96

snapshot repository volume, 95, 191

Snapshot Volume

- description of, 191
- snapshot repository volume, 191

snapshot volume

- failure due to role reversal, 98
- with Remote Volume Mirroring feature, 95

SNMP destinations

- specifying, 76
- where stored, 11

SNMP trap message

- interpreting, 76
- use for critical alerts, 2

software components, 8

source volume

- properties of a, 207
 - selecting a, 199
 - source volumes, 184
 - specific volume-to-LUN mapping, 51
 - start percentage, cache flush, 34
 - status icon
 - changes during a recovery procedure, 68
 - changes during recovery procedure, 67
 - Stop Copy option, 212
 - stop percentage, cache flush, 34
 - Storage Array
 - comment, 73
 - configuring, 25
 - connectivity, 5
 - defined, 2
 - directly managed, 7
 - host-agent managed, 7
 - indications of failure, 63
 - locating, 77
 - logical components, 3
 - management activities, 5
 - physical components, 3
 - renaming, 77
 - synchronizing clocks, 78
 - tuning options available, 70
 - viewing properties, 78
 - storage array component failures, 176
 - storage array health, summary of conditions, 55
 - Storage Array Management, 1
 - storage management software
 - where installed, 1
 - storage management software requirements, 101
 - storage management station
 - defined, 2
 - described, 1
 - Storage Partition
 - creating, 39, 74
 - defined, 4
 - major steps to creating, 49
 - Topology, 50
 - Storage Partitioning, 189
 - overview and example, 47
 - topology, 44
 - Wizard, 47
 - storage partitioning
 - and secondary volumes, 94
 - storage partitioning, description of, 189
 - switch zoning
 - and switch ports, 114, 121, 127
 - switches
 - cascading, 109
 - list of certified, 101
 - zoning, 109
 - Synchronization Priority Level
 - defined, 147
 - setting the, 89, 143
 - Synchronization-in-Progress status, 87, 88, 90
 - Synchronizing status, 180
 - synchronous mirroring, 84
- ## T
- target volume
 - creating a, 197
 - disable read-only attribute, 216
 - enable read-only attribute, 217
 - properties of a, 207
 - recommended capacity of, 197
 - selecting a, 197
 - setting read and write permissions, 216
 - target volumes, 184
 - task reference, 73
 - terminology
 - functional equivalents, vi
 - text formatting conventions, vi
 - topological elements, defined, 44
 - Topology view, 44

transfer rate, 70
troubleshooting problems, 242

U

unassigned capacity
 see unconfigured capacity
unconfigured capacity
 adding to a storage array, 74
 creating volumes from, 74
 creating volumes on, 42
 defined, 4
Unconfigured Capacity Node, 4
Understanding guide
 how to use, v
 organization, v
 prerequisite knowledge, v
Unresponsive status
 defined, 58
 recovery, 76
Unsynchronized status, 180
 after link interruption, 86
upgrading mirror repository volumes, 139
uplink port, 114, 121
UPS, 34
Utilities
 hot_add, 223
 SMrepassist, 205

V

view
 drives, 200
View Drives dialog, 200
volume
 and I/O accessibility, 80
 candidates for mirroring, 83

 creating, 39
 defined, 3
 deleting primary or secondary, 155
 initializing, 77
 missing, 171, 243
 number of mirror relationships, 80
 renaming, 77
 secondary, 191
 sharing access to a, 51
 snapshot, 191
 snapshot repository, 191
 viewing properties, 78
 volume-to-LUN settings, 50
volume copies, maximum number allowable, 194
Volume Copy
 basic concepts, v, 183
 change copy priority, 214
 copy priority, 187
 copy status, 209
 critical events, 246
 Feature Key File for, 195
 invalid source volumes, 200
 invalid target volumes, 197
 maximum number of copies, 186
 monitoring progress of, 208
 premium feature, enabling the, 195
 progress of a, 209
 re-copying a, 210
 remove copy pairs, 213
 restrictions, 194
 stopping a, 212
 troubleshooting, 242
 usage scenarios, 218
 valid source volumes, 200
 valid target volumes, 197
 viewing properties of, 207
volume group
 capacity of, 30
 data stripe, 71
 defined, 3

- deleting, 74
- initializing, 77
- locating, 77
- moving to another controller, 71
- redistributing, 77

volume modification priority
 see modification priority

volume-to-LUN mapping, 189

- default, 51
- defined, 46, 50
- preference, 42
- shared, 51
- specific, 51

W

World Wide Name (WWN)

- for a source or target volume, 207
- updates, 173

write cache mirroring

- described, 34
- how to enable, 34

write caching

- and data loss, 34
- and performance, 34
- described, 72
- without batteries, 34

write mode, changing, 154

WWN (World Wide Name), 45