

SGI® TP9400 and SGI® TP9500 Software Concepts Guide

007-4305-003

CONTRIBUTORS

Written by Jenn Byrnes

Illustrated by Chrystie Danzer

Edited by Susan Wilkening

Production by Glen Traefald

Engineering contributions by Bill Andrews, Marty Castilla, and Robert Novak.

COPYRIGHT

© 2002 Silicon Graphics, Inc. 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 Silicon Graphics, Inc.

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., 1600 Amphitheatre Pkwy 2E, Mountain View, CA 94043-1351.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, and IRIX are registered trademarks of Silicon Graphics, Inc.

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.

Cover design by Sarah Bolles, Sarah Bolles Design, and Dany Galgani, SGI Technical Publications.

Record of Revision

Version	Description
001	August 2000 Original publication.
002	January 2001 Engineering revisions.
003	October 2002 Updated to support SGI TP9400 and SGI TP9500.

Contents

Record of Revision	iii
Figures	xi
Tables	xiii
About This Guide.	xv
Software and Hardware Terminology	xv
Related Publications	xvi
Obtaining Publications	xvii
Conventions	xviii
Reader Comments	xviii
1. Storage Array Management Architecture and Software Components.	1
Major Hardware Components	1
Physical Components of a Storage Array	3
Logical Components of a Storage Array	4
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	10
Device Tree	12
Device Table	13
Enterprise Management Window Menus	14
Enterprise Management Window Toolbar	15

Array Management Window 16
Array Management Window Tabs 17
Array Management Window Menu Options 18
Array Management Window Toolbar Options 20
2. Configuring Storage Arrays 23
Protecting Your Data 23
Software Redundancy 23
Use Remote Volume Mirroring 24
Create a Snapcopy Volume 24
RAID Levels and Data Redundancy 24
Hardware Redundancy 27
Protect Data in the Controller Cache Memory 27
Configure Hot Spare Drives 29
Enable a Background Media Scan 29
Use Available Channel Protection When Possible 29
I/O Data Path Protection 30
Enable Multi-Path Driver with AVT 30
Disable Multi-Path Drive with AVT 31
Storage Array Password Protection 31
Creating Volumes and Storage Partitions 31
Configuring Storage Array Capacity into Volumes 32
Register the Volume with the Operating System 33
Dynamic Volume Expansion (DVE) 33
Mappings View 34
Storage Partitioning 38
Storage Partitioning Example 39
Major Steps to Create Storage Partitions 40
Heterogeneous Hosts 43
Heterogeneous Hosts Example 44
Maintaining Storage Arrays in the Management Domain 46
Monitoring the Health of Storage Arrays 46
Storage Array Status Icons 48

Event Monitor	49
Alert Notifications	52
Customer Support Alert Notifications	53
Problem Notification	54
Recovering from Storage Array Problems	56
Recovery Guru Example	57
Tuning Storage Arrays	61
Task Reference	64
3. Remote Volume Mirroring	71
Feature Concepts and Requirements	71
Introduction to Remote Volume Mirroring	71
Primary and Secondary Volumes	72
Mirror Repository Volumes.	73
Mirror Relationships.	74
Data Replication	75
Link Interruptions or Secondary Volume Errors	76
Connectivity and Input/Output	78
Fibre Channel Fabric Configurations	78
Volume Ownership	79
Distance Between Primary and Secondary Storage Array	80
General Performance Considerations	81
Using Other Premium Features and Remote Volume Mirroring	82
Storage Partitioning	82
Snapcopy Volumes	83
Hardware and Software Requirements.	86
Hardware Requirements	87
Software Requirements	87
Recommended Configurations	88
Highest Availability Campus Configuration	88
Campus Configuration	90
Intra-Site Configuration.	91

Hardware and Software Checklist 92
Hardware Checklist 92
Software Checklist 93
Hardware and Software Installation 94
Pre-Installation Requirements 94
Site Preparation 94
Hardware and Software Requirements. 94
Switch Zoning Overview 95
Hardware Installation. 97
Highest Availability Campus Configuration 98
Campus Configuration	107
Intra-Site Configuration	114
Software Installation	121
Verifying Installation and Configuration	122
Using the Remote Volume Mirroring Premium Feature	123
Enabling and Activating the Remote Volume Mirroring Feature	124
Determining Remote Volume Mirroring Feature Status	124
Enabling the Remote Volume Mirroring Feature	126
Activating the Remote Volume Mirroring Feature	126
Creating Mirror Relationships	128
Creating Volumes for Mirror Relationships	129
Creating a Mirror Relationship Using the Create Remote Volume Mirror Wizard	129
Maintaining Mirror Relationships	131
Viewing Mirror Relationship Details	132
Changing the Synchronization Priority Level	133
Removing Mirror Relationships.	134
Deleting Primary or Secondary Volumes	135
Deactivating and Disabling the Remote Volume Mirroring Feature.	136
Deactivating the Remote Volume Mirroring Feature.	136
Disabling the Remote Volume Mirroring Feature.	136

Disaster Recovery137
Planning for Disaster Recovery137
Backing Up Critical Data137
Preparing the Secondary Storage Array138
Ensuring Failover Protection Using Host Clustering Software138
Disaster Preparation Checklist139
Recovering from Switch Failures139
Recovering from Storage System Failures140
Entire Site Failover Procedure141
Rebuilding a Damaged Site142
Re-creating a Mirror Relationship142
Reversing the Roles of Primary and Secondary Volumes143
Changing a Secondary Volume to a Primary Volume144
Changing a Primary Volume to a Secondary Volume144
Troubleshooting145
General Troubleshooting.146
Inappropriate Fibre Channel Fabric Configurations153
Storage Array Component Failures.154
General Switch Troubleshooting155
Remote Volume Mirroring Critical Events.156
Controller Firmware Version 5.30 Upgrades158
Index161

Figures

Figure 1-1	Physical Components of a Storage Array	3
Figure 1-2	I/O Data Connection	6
Figure 1-3	Network Management Connections	6
Figure 1-4	Enterprise Management Window	9
Figure 1-5	Array Management Window	10
Figure 1-6	Enterprise Management Window	12
Figure 1-7	Device Tree with a Management Domain	13
Figure 1-8	Array Management Window	17
Figure 2-1	Mappings View Window	35
Figure 2-2	Storage Partitioning Example	40
Figure 2-3	Defining Hosts and Host Groups in the Topology View	43
Figure 2-4	Host Port Definitions Dialog	44
Figure 2-5	Heterogeneous Hosts Example	45
Figure 2-6	Monitoring the Health of Storage Arrays Using the Enterprise Management Window.	47
Figure 2-7	Event Monitor Configuration	50
Figure 2-8	Event Monitor Example	51
Figure 2-9	Problem Notification in the Array Management Window.	55
Figure 2-10	Recovery Guru Toolbar Button	56
Figure 2-11	Recovery Guru Window	58
Figure 2-12	Status Changes During Example Recovery Operation	59
Figure 2-13	Status Changes When an Example Recovery Operation Completes	60
Figure 3-1	Primary and Secondary Volumes Displayed in the Array Management Window.	73
Figure 3-2	Mirror Repository Volumes Displayed in the Array Management Window	74
Figure 3-3	Data Replication between Storage Arrays	76

Figure 3-4	Mirrored Volume Pair in Unsynchronized Status77
Figure 3-5	Mirrored Volume Pair in Synchronization-in-Progress Status77
Figure 3-6	Controller A as Controller Owner.79
Figure 3-7	Controller B as Controller Owner80
Figure 3-8	Storage Partitioning Example83
Figure 3-9	Snapcopy Volume Included in a Mirror Relationship84
Figure 3-10	Snapcopy Volume Associated with Primary Volume Accounting85
Figure 3-11	Snapcopy Volume Failure Because of Role Reversal86
Figure 3-12	Highest Availability Campus Configuration.89
Figure 3-13	Campus Configuration90
Figure 3-14	Intra-Site Configuration91
Figure 3-15	Switch Zoning in Cascaded and Non-Cascaded Fibre Channel Switches96
Figure 3-16	Highest Availability Campus Configuration.99
Figure 3-17	Switch Zoning for the Highest Availability Campus Configuration	101
Figure 3-18	Host Bus Adapter Connections to Fibre Channel Switches.	105
Figure 3-19	Storage Array Connections to Fibre Channel Switches	106
Figure 3-20	Connecting Switches over a Remote Network to Complete Fabric Environments	107
Figure 3-21	Campus Configurations	108
Figure 3-22	Switch Zoning for the Campus Configuration	109
Figure 3-23	Host Bus Adapter Connections to Fibre Channel Switches.	112
Figure 3-24	Storage Array Connections to Fibre Channel Switches	113
Figure 3-25	Connecting Switches over a Remote Network to Complete the Fabric Environment.	113
Figure 3-26	Intra-Site Configuration	115
Figure 3-27	Switch Zoning for the Intra-Site Configuration	116
Figure 3-28	Primary Host Bus Adapter Connections to Fibre Channel Switches	119
Figure 3-29	Secondary Host Bus Adapter Connections to Fibre Channel Switches	119
Figure 3-30	Primary Storage Array Connections to Fibre Channel Switches	120
Figure 3-31	Secondary Storage Array Connections to Fibre Channel Switches	121
Figure 3-32	Remote Volume Mirroring Feature Status	124
Figure 3-33	Synchronizing (Left) and Optimal (Right) Mirror Relationship	131

Tables

Table i	Terminology Change Reference	xvi
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	4
Table 1-4	Major Software Components	8
Table 1-5	Enterprise Management Window Menus	14
Table 1-6	Enterprise Management Window Toolbar Buttons	15
Table 1-7	Array Management Window Tabs	18
Table 1-8	Array Management Window Menus	19
Table 1-9	Array Management Window Toolbar Buttons	20
Table 2-1	RAID Level Configurations	26
Table 2-2	Creating Volumes from Unconfigured or Free Capacity	32
Table 2-3	Mappings View Tab	36
Table 2-4	Volume-to-LUN Terminology	36
Table 2-5	Storage Array Status Indications	48
Table 2-6	Tuning Options in Array Management Window	61
Table 2-7	Task Reference: Populating a Management Domain	64
Table 2-8	Task Reference: Storage Array Configuration	65
Table 2-9	Task Reference: Performance and Tuning	66
Table 2-10	Task Reference: Event Notification	67
Table 2-11	Task Reference: Recovering From Problems	67
Table 2-12	Task Reference: Miscellaneous System Administration	68
Table 3-1	Hardware Checklist	92
Table 3-2	Software Checklist	93
Table 3-3	Highest Availability Campus Configuration Connections	101
Table 3-4	Campus Configuration Connections	110

Table 3-5	Intra-Site Configuration Connections	116
Table 3-6	Disaster Preparation Checklist	139

About This Guide

This guide explains the terminology and features of the Total Performance Storage System Manager (TPSSM) software for SGI TP9400 and SGI TP9500. This guide is written primarily for the system administrator.

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, “Storage Array Management Architecture and Software Components” describes the major hardware components in the TPSSM 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.

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. See Table i

if you are upgrading storage array hardware and migrating from earlier storage management software versions.

Table i Terminology Change Reference

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 TP9400 RAID Installation and Upgrade Guide* This guide gives complete instructions on how to unpack, install, and configure the SGI TP9400 and its components. It also contains upgrade information.
- *SGI TP9500 RAID Installation and Upgrade Guide* This guide gives complete instructions on how to unpack, install, and configure the SGI TP9500 and its components. It also contains upgrade information.

- *SGI TP9400 and SGI TP9500 RAID Owner's Guide* for controller and drive enclosure hardware. This guide provides installation and site requirements, installation instructions, product overviews, and information about component routine maintenance, operation, troubleshooting, and replacement.
- *SGI TP9400 and SGI TP9500 RAID Administration Guide* for the storage management software specific to the IRIX 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 the IRIX operating environment.

When you have completed the tasks in the *SGI TP9400 and SGI TP9500 RAID 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 TP9400 and SGI TP9500 RAID Administration Guide* for information specific to the IRIX 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, please tell us. 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 us in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the webpage of the Technical Publications Library at:
<http://docs.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:
Technical Publications
SGI
1600 Amphitheatre Pkwy., M/S 535
Mountain View, California 94043-1351
- Send a fax to the attention of “Technical Publications” at +1 650 932 0801.

We value your comments and will respond to them promptly.

Storage Array Management Architecture and Software Components

This chapter describes the major hardware components in the Total Performance Storage System Manager (TPSSM) 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 the host or a GUI workstation. The storage commands will be sent over the host I/O data connection (when connected in-band) or through TCP/IP over the network (when connected directly).

Major Hardware Components

Table 1-1 lists the major hardware components that are involved in the TPSSM storage management architecture. For a detailed description of the network setup requirements,

see the appropriate hardware site preparation and installation 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). See “Physical Components of a Storage Array” on page 3 and “Logical Components of a Storage Array” on page 4 for more information.
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 shown in Figure 1-1 and described in Table 1-2. In some newer hardware, the configuration or location might be slightly modified, but the type and use of the components remains the same.

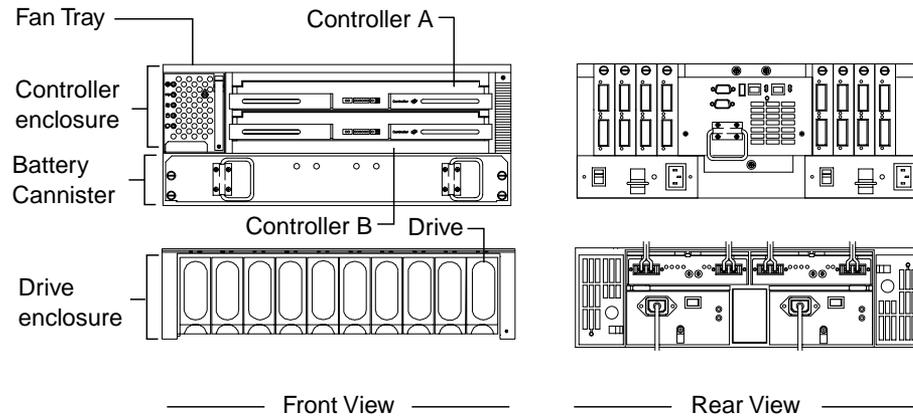


Figure 1-1 Physical Components of a Storage Array

Table 1-2 describes the physical components shown in Figure 1-1.

Table 1-2 Physical Components of a Storage Array

Component	Description
Controller	Board and firmware that controls 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.
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 1-8 on page 17). 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 38.

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 38).
- 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 1-2.

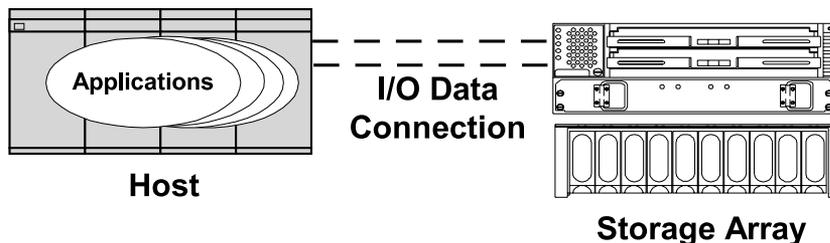


Figure 1-2 I/O Data Connection

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 1-3, is called the *network management connection*. Two kinds of network management connections are described in “Types of Network Management Connections” on page 6.

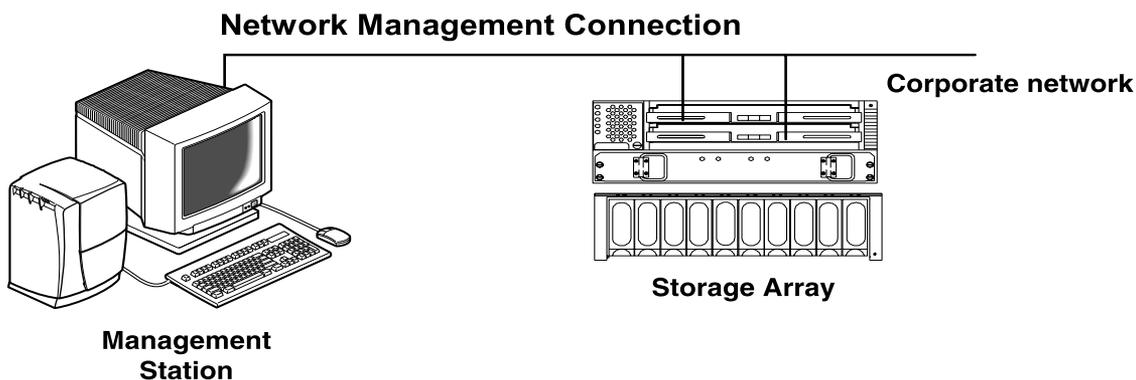


Figure 1-3 Network Management Connections

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 TP9400 and SGI TP9500 RAID 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 TPSSM 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 "Related Publications" on page xvi.

For more information about host-agent managed storage arrays, see the *SGI TP9400 and SGI TP9500 RAID Administration Guide*.

Major Software Components

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

Table 1-4 Major Software Components

Component	Description
TPSSM 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 TP9400 and SGI TP9500 RAID Administration Guide</i> for details.</p> <p>For more information, see “Client Software Windows” on page 9.</p>
TPSSM 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 TP9400 and SGI TP9500 RAID Administration Guide</i> for details.</p>
TPSSM Util	<p>The <code>util</code> package software contains the <code>tpssmdevices</code> utility that correlates the volumes you create using the storage management software with their operating system device names. See the <i>SGI TP9400 and SGI TP9500 RAID Administration Guide</i> for details.</p>

Client Software Windows

The client software has two main windows: the **Enterprise Management** window, shown in Figure 1-4, and the **Array Management** window, shown in Figure 1-5.

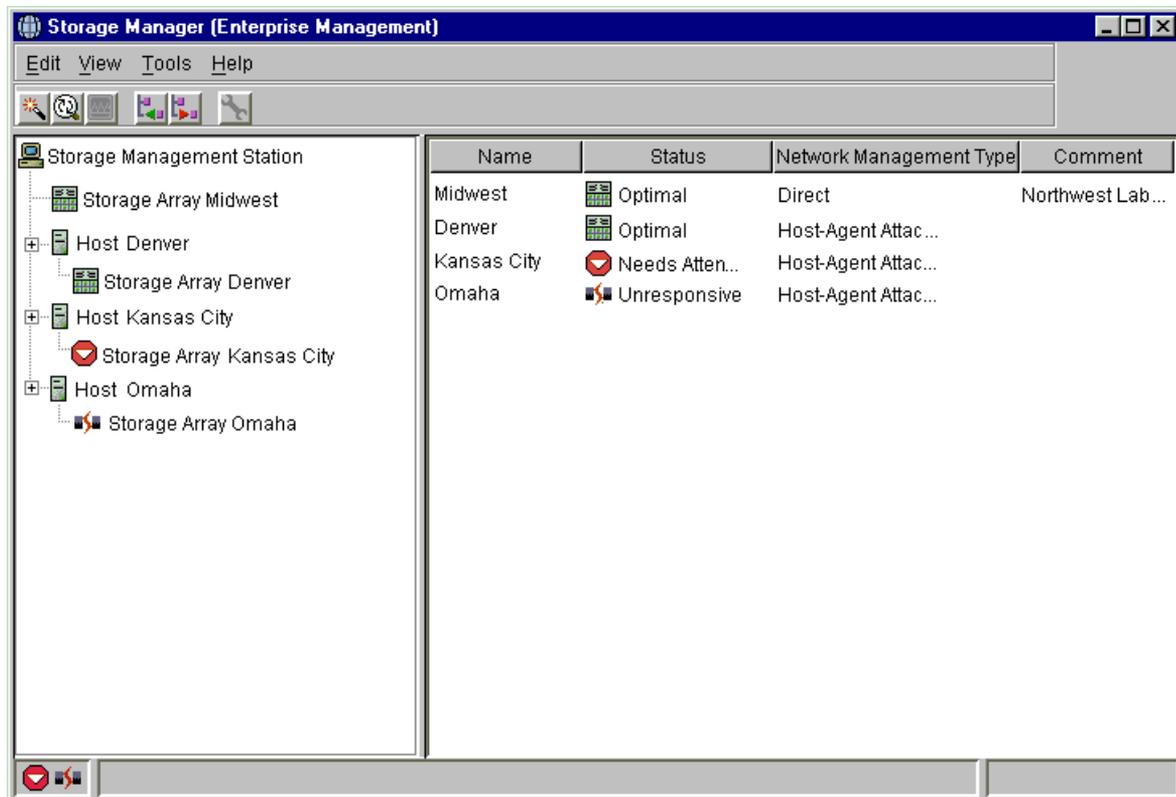


Figure 1-4 Enterprise Management Window

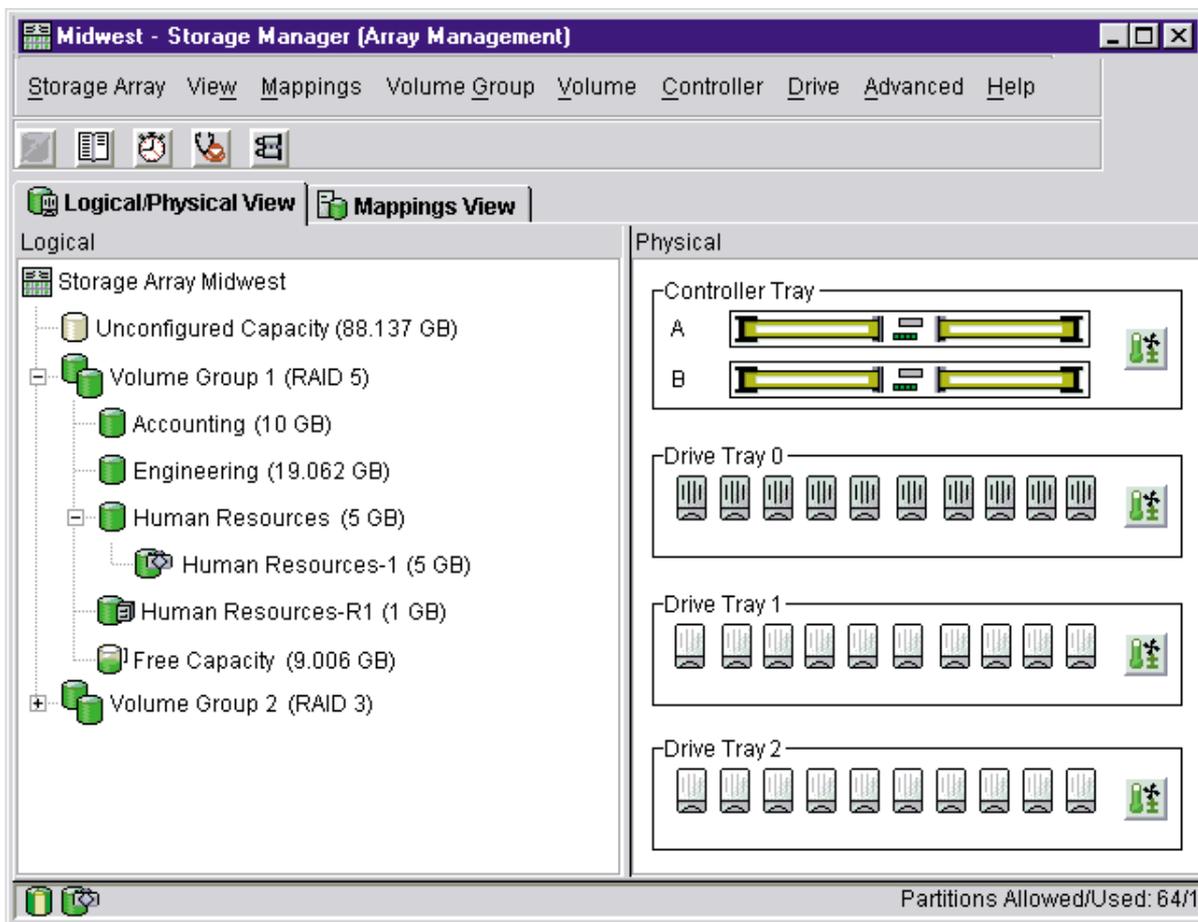


Figure 1-5 Array Management Window

Enterprise Management Window

The **Enterprise Management** window, shown in Figure 1-6, 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 xvi).

- 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 46.

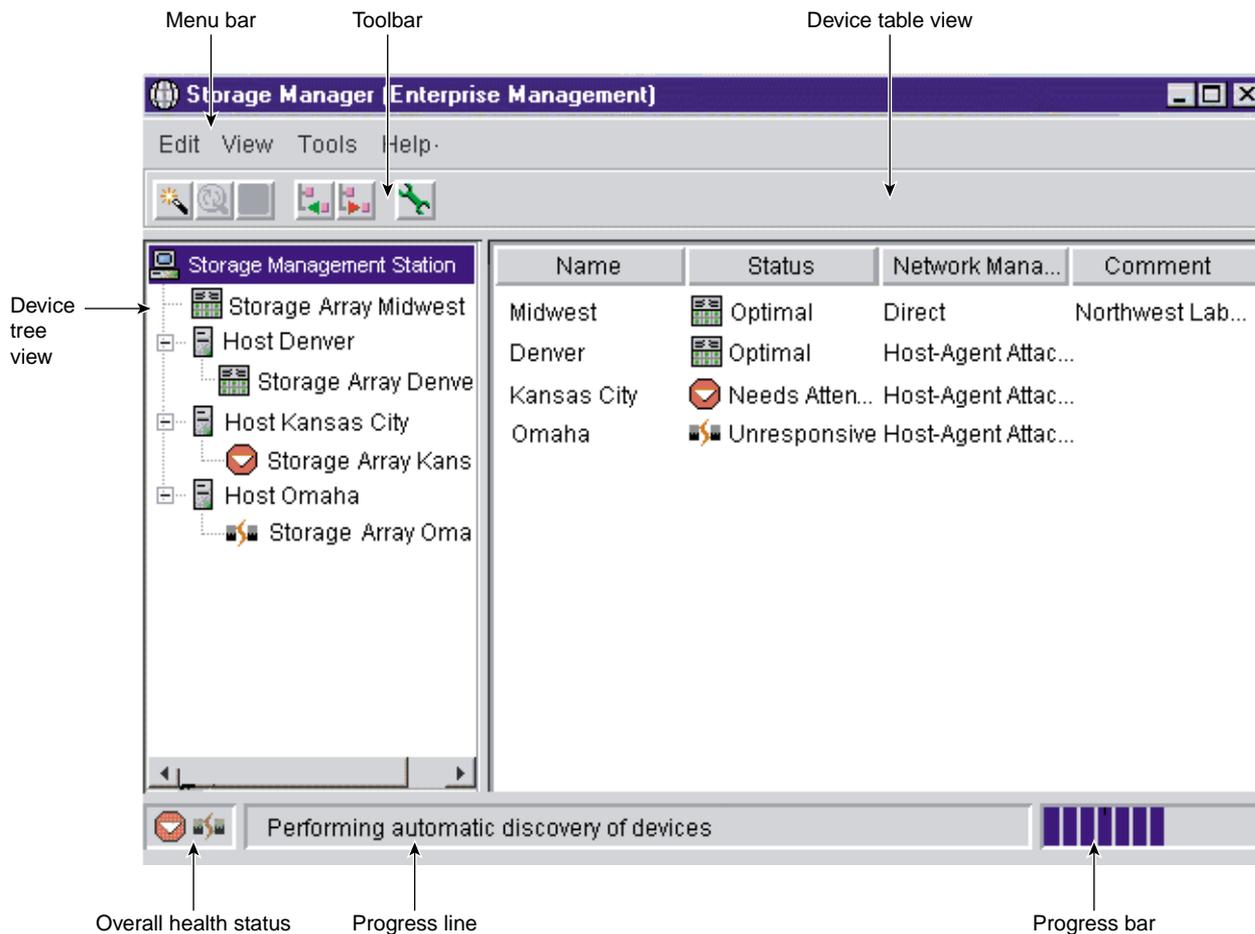


Figure 1-6 Enterprise Management Window

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 1-7 shows an example Device Tree.

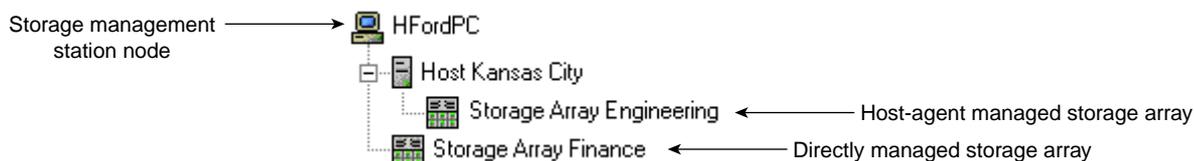


Figure 1-7 Device Tree with a Management Domain

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.

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 1-6 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	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. 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).
View	Allows you to sort the entries in the Device Table by Name, Status, Management Type, or Comment. 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. For more information on partially managed storage arrays, see the Enterprise Management window Help.
Tools	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.
Help	Allows you to view the Help system and software version and copyright information for the Enterprise Management software.

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.
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.
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 1-6 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 1-8 on page 17.

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.

The storage management software supports firmware version 5.30 and all firmware versions 4.x and 5.x. For maximum system stability, the recommended minimum is firmware version 4.01.02.30. However, to access all the features of version 8.30, you must upgrade to firmware version 5.30.

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 5.30 and the other has firmware version 4.x. 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 5.30 will use version 8.30 of the storage management software while the storage array with firmware version 4.x will use version 7.x. 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.

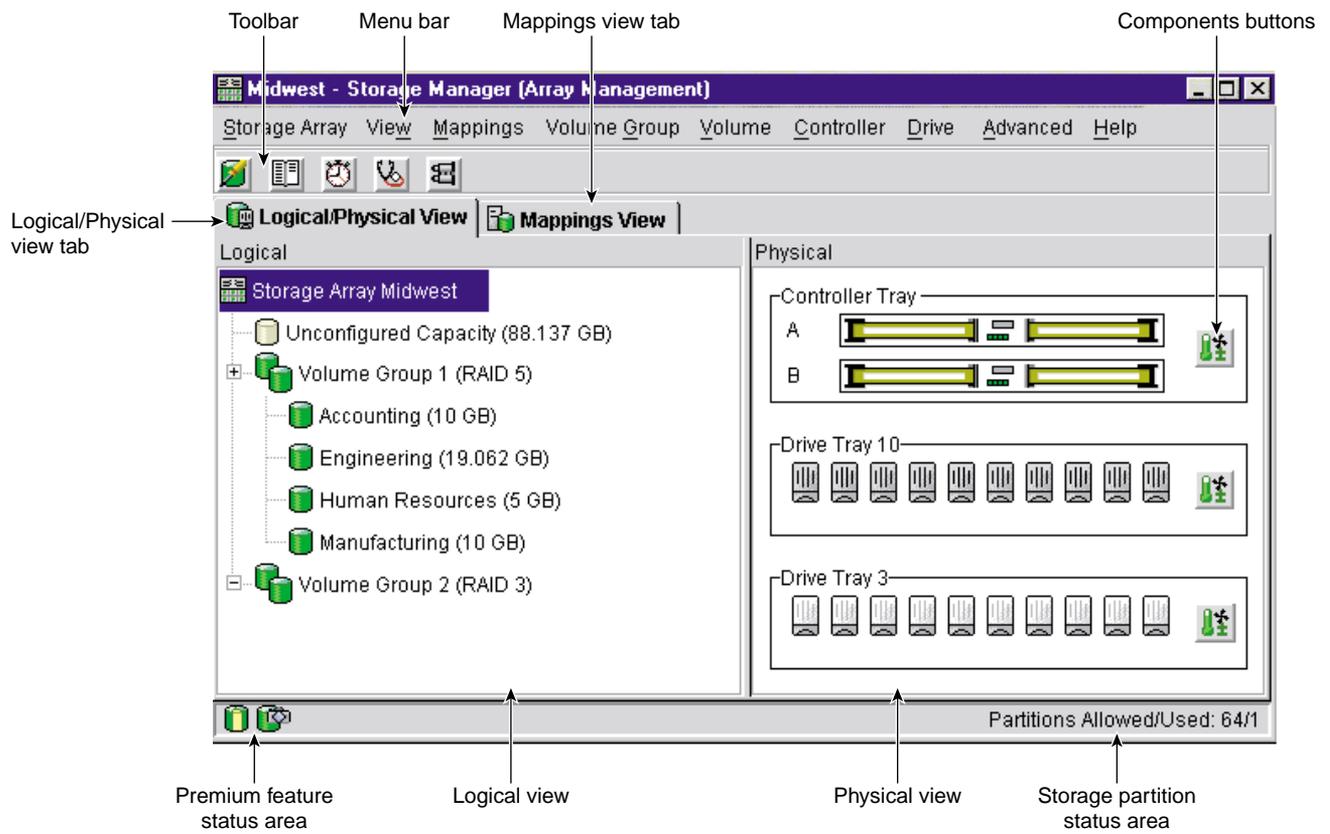


Figure 1-8 Array Management Window

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 1-8 on page 17. 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 1-8 on page 17 shows the organization of storage array capacity into volume groups and volumes.</p> <p>The Physical view in the right pane of Figure 1-8 on page 17 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 34.</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 19. The menus are used to perform storage management operations for a selected

storage array or for selected components within a storage array. See Figure 1-8 on page 17 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 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.
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.

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

Menu	Description
Advanced	Presents maintenance options, which should only be used under the guidance of technical support.
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 1-8 on page 17 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.

Table 1-9 Array Management Window Toolbar Buttons (continued)

Toolbar button	Description
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.

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 23
- “RAID Levels and Data Redundancy” on page 24
- “Hardware Redundancy” on page 27
- “I/O Data Path Protection” on page 30
- “Storage Array Password Protection” on page 31

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 Snapcopy Volume Creation 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 Snapcopy Volume

A snapcopy 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 snapcopy volume is created so that an application, such as a backup application, can access the snapcopy volume and read the data while the base volume remains online and user-accessible. When the backup completes, the snapcopy volume is no longer needed. You can also create several snapcopy volumes of a base volume and write data to the snapcopy volumes to perform testing and analysis. For example, before upgrading a database management system, snapcopy 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, and 5. Table 2-1 on page 26 describes these configurations. Each level provides different performance and protection features.

RAID levels 1, 3, and 5 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.

Note: 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.

Table 2-1 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	Striping/Mirroring Mode	<p>RAID 1 is also called RAID 10 or 0+1. It offers high performance and uses disk mirroring to make an exact copy from one drive to another drive.</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>

Table 2-1 RAID Level Configurations (**continued**)

RAID Level	Short Description	Detailed Description
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>
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>

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: Sometimes write caching is disabled when batteries are low or discharged. If you enable a parameter called Write Caching Without Batteries on a volume, write caching continues even when batteries in the controller enclosure or array enclosure are

discharged. If you do not have a Universal Power Supply (UPS) for power protection, do not enable this parameter because data in the cache will be lost during a power outage when the controller enclosure or array enclosure does not have working batteries.

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 2-6 on page 61), and therefore the I/O request rate (also described in Table 2-6 on page 61). 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 TPSSM 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.

For more information about controller enclosure batteries or the UPS, see the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.

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, or 5 volume. The hot spare drive adds another level of redundancy 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, the data from the hot spare is copied back to the replacement drive. This is called *copyback*.

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, or 5 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*.

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.

Use Available Channel Protection When Possible

In a SCSI environment, whether or not the volume group candidate has channel protection depends on the RAID level of the volume group and how many drives are

present on any single drive channel. For example, a RAID 5 volume group would not have channel protection if more than one drive was present on a single drive channel.

In a Fibre Channel environment, channel protection is usually present for any volume group candidate because when the storage array is properly cabled there are two redundant arbitrated loops for each drive.

I/O Data Path Protection

I/O data path protection to redundant controllers in a storage array is accomplished with the Auto-Volume Transfer (AVT) feature and a host multi-path driver.

A multi-path driver is an I/O path failover driver, such as RDAC (Redundant Disk Array Controller), that is installed on host computers and accesses the storage array. AVT is a feature of the controller firmware.

AVT is disabled by default, but if the host operating system supports it, AVT will be automatically enabled based on the failover options supported by the specified host operating system.

For operating system-specific failover options, see the *SGI TP9500 RAID Installation and Upgrade Guide*.

Enable Multi-Path Driver with AVT

If AVT is enabled when a volume is created, a controller must be assigned to own the volume (called the preferred controller, or preferred owner). 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 to fail, the multi-path driver will issue the I/O to the alternate controller.

When AVT is enabled and used in conjunction with a host multi-path driver, it helps ensure an I/O data path is available for the storage array volumes. The AVT feature changes the ownership of the volume receiving the I/O to the alternate controller.

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

Disable Multi-Path Drive with AVT

When AVT is disabled, the I/O data path will still be 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 an individual volume, and a problem occurs along the data path to its preferred controller, **all** volumes on that controller will be transferred to an alternate controller instead of just that particular volume.

Storage Array Password Protection

Caution: Executing destructive commands on a storage array can cause serious damage, including data loss. Without password protection, **all** options are available within this storage management software.

For added security, you can configure a password for each storage array you manage. Because the password is stored on the storage array, each storage array that you want to be password protected will need a password. A specified password protects any options that the controller firmware deems destructive. These options include any functions that change the state of the storage array such as creation of volumes, modification of cache settings, and so on.

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.

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

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.

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 2-2.

Table 2-2 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> <p>For a description of the Volume-to-LUN Mapping Preference parameter see “Related Publications” on page xvi.</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>

Register the Volume with the Operating System

Note: The `hot_add` utility is not available for all operating systems. See the *SGI TP9500 RAID Installation and Upgrade Guide* for information on operating system availability and how to run this utility.

After creating all volumes and assigning volume-to-LUN mappings, the host-based `hot_add` utility is used to register the volume with the operating system. The host-based `tpssmdevices` utility (if applicable for your operating system) is used to associate the mapping between the physical device name and the volume name. See “Storage Partitioning” on page 38 for more information.

Dynamic Volume Expansion (DVE)

Caution: Increasing the capacity of a standard volume is only supported on certain operating systems. If volume capacity is increased on a host operating system that is unsupported, the expanded capacity will be unusable, and you cannot restore the original volume capacity. For information on supported operating systems, see the *TP9500 Release Notes* shipped with the storage management software.

Volume capacity of standard volumes and snapcopy repository volumes is increased by using Dynamic Volume Expansion (DVE). This feature expands the capacity of an existing volume by either using a volume’s free capacity on an existing volume group, or by adding unconfigured capacity (new or unassigned drives) to the volume group. You can expand a volume dynamically without losing access to it or to any other volumes.

Increasing the capacity of a snapshot repository volume does not increase the capacity of the associated snapcopy volume. However, unless the capacity of the base volume is also increased, it is not necessary to increase the size of the snapcopy volume. The snapcopy volume’s capacity is always based on the capacity of the base volume at the time the snapcopy volume is created.

You cannot add storage capacity to a volume if:

- One or more hot spare drives are in use in the volume.
- The volume has a non-optimal status.
- Any volume in the volume group is in any state of modification.

- The controller that owns this volume is in the process of adding capacity to another volume. (Each controller can add capacity to only one volume at a time.)
- No free capacity exists in the volume group, **and** no unconfigured capacity is available to add to the volume group.
- No unconfigured capacity (in the form of drives) is available to add to the volume group.

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 2-3.

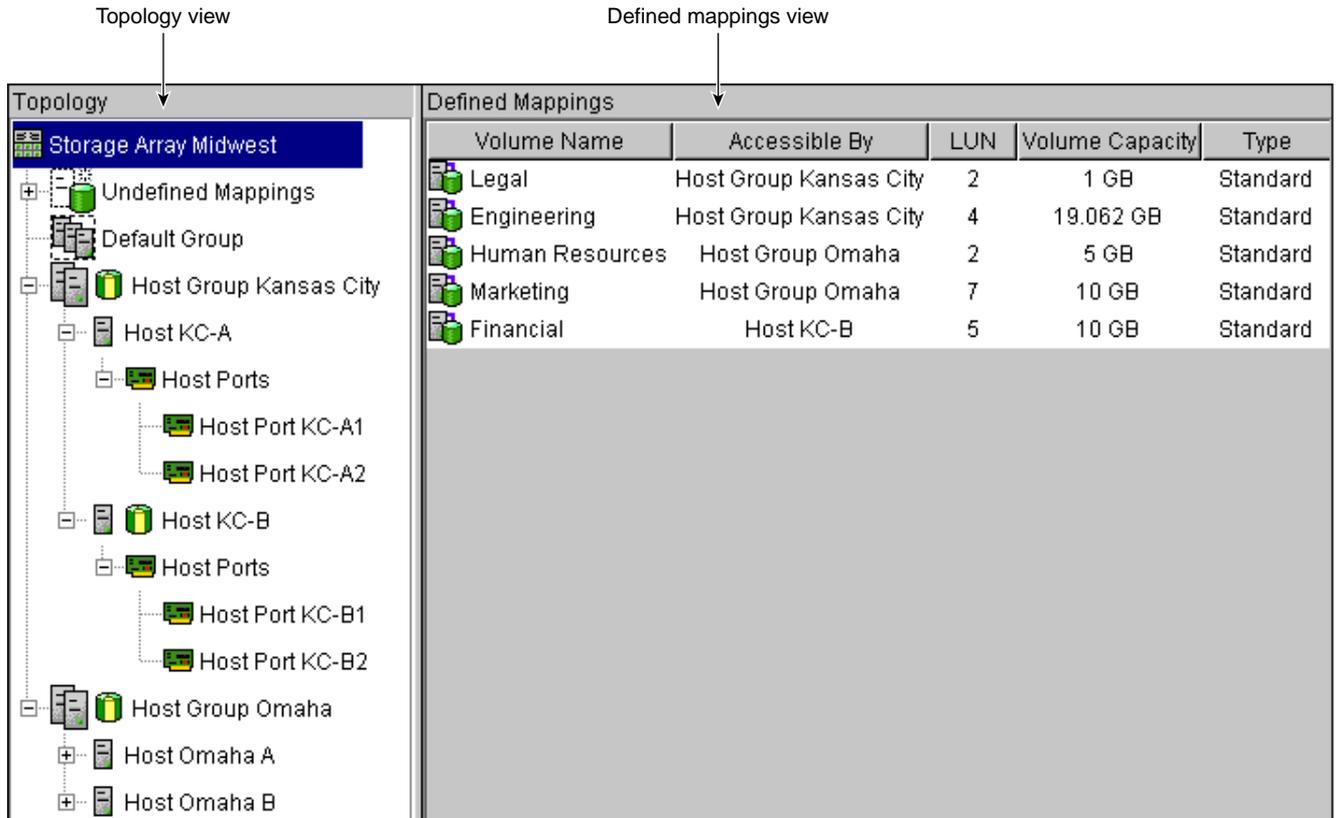


Figure 2-1 Mappings View Window

Table 2-3 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 2-4 on page 36 describes the topological elements displayed in the **Mappings View** window in Figure 2-1 on page 35.

Table 2-4 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.

Table 2-4 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 45 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 2-4 Volume-to-LUN Terminology (**continued**)

Term	Description
Default Volume-to-LUN mapping	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 .
Specific Volume-to-LUN mapping	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.

Storage Partitioning

Storage Partitioning is a premium feature of the storage management software and must be enabled either by you or your storage vendor.

A storage partition is a logical entity consisting of one or more storage array volumes that can be shared among hosts, which may be part of a host group or accessed by a single host. A storage partition is created when you define a single host or a collection of hosts, called a host group, and then define a volume-to-LUN mapping. This mapping defines what host or host group will have access to a particular volume in your storage array. Hosts and host groups can only access data through assigned volume-to-LUN mappings. Each volume can be mapped to one storage partition.

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.

A maximum of 64 storage partitions can be supported by the storage management software. The 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. 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 (heterogeneous hosts) 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.

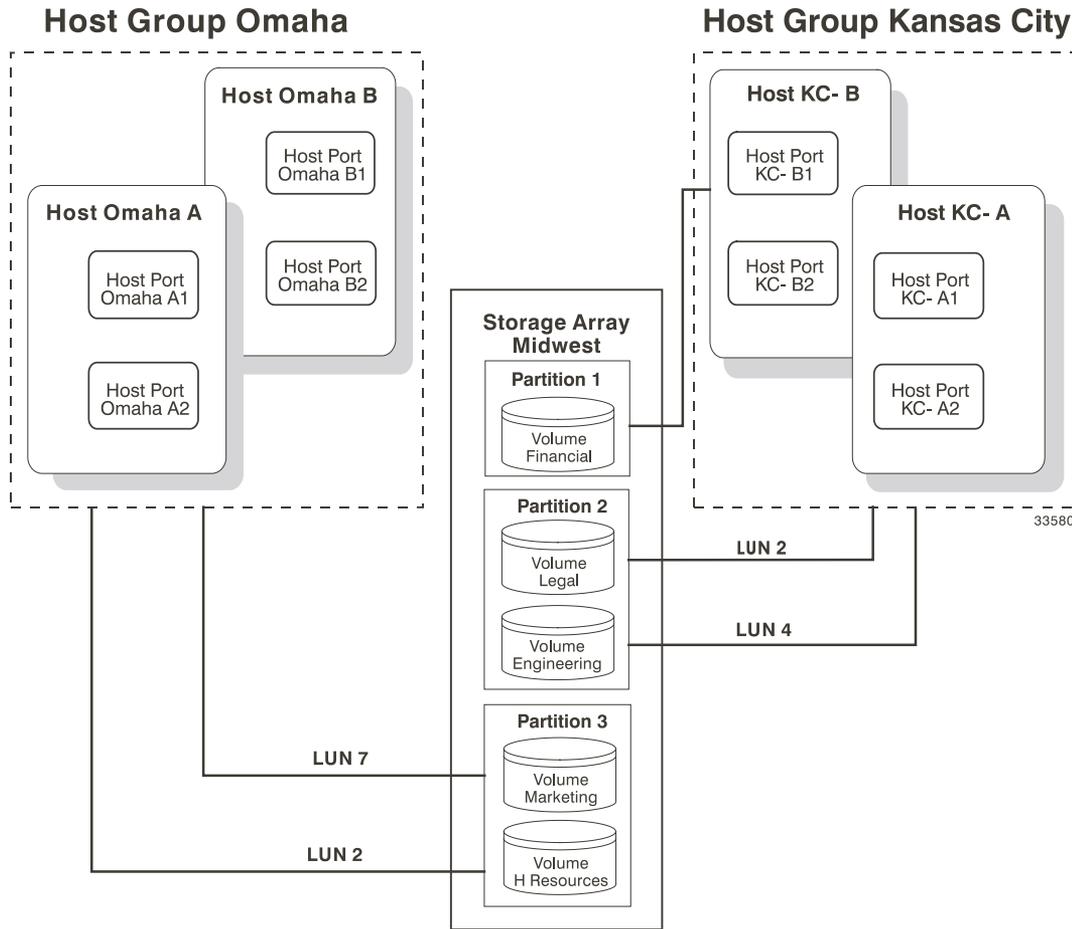


Figure 2-2 Storage Partitioning Example

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).

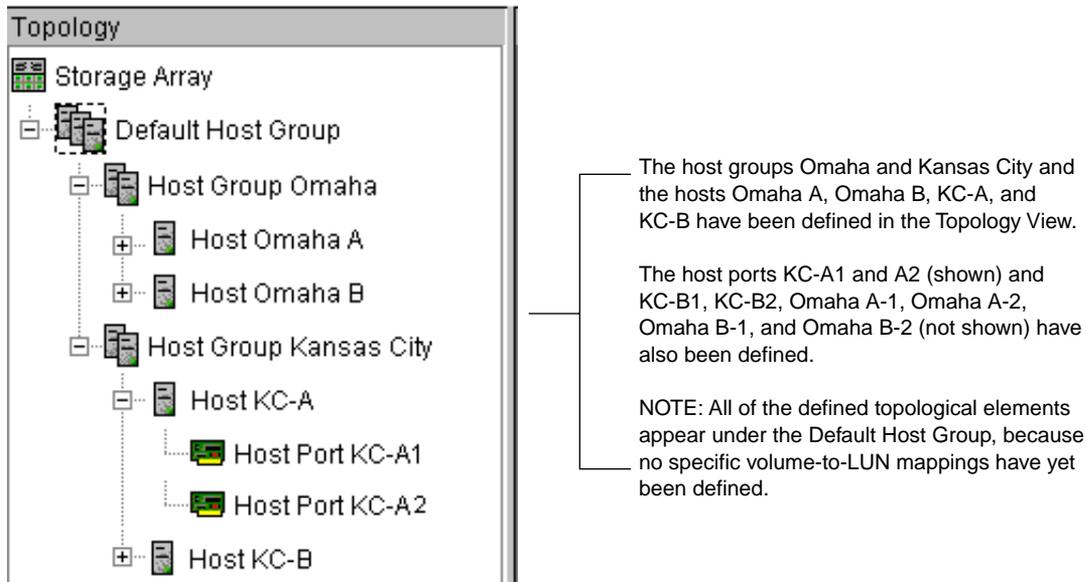


Figure 2-3 Defining Hosts and Host Groups in the **Topology** View

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 45, 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.

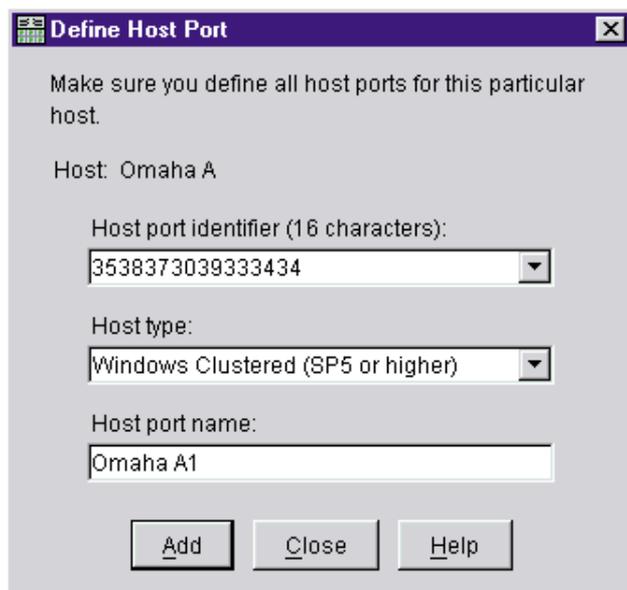


Figure 2-4 Host Port Definitions Dialog

In the example shown in Figure 2-5 on page 45, 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.

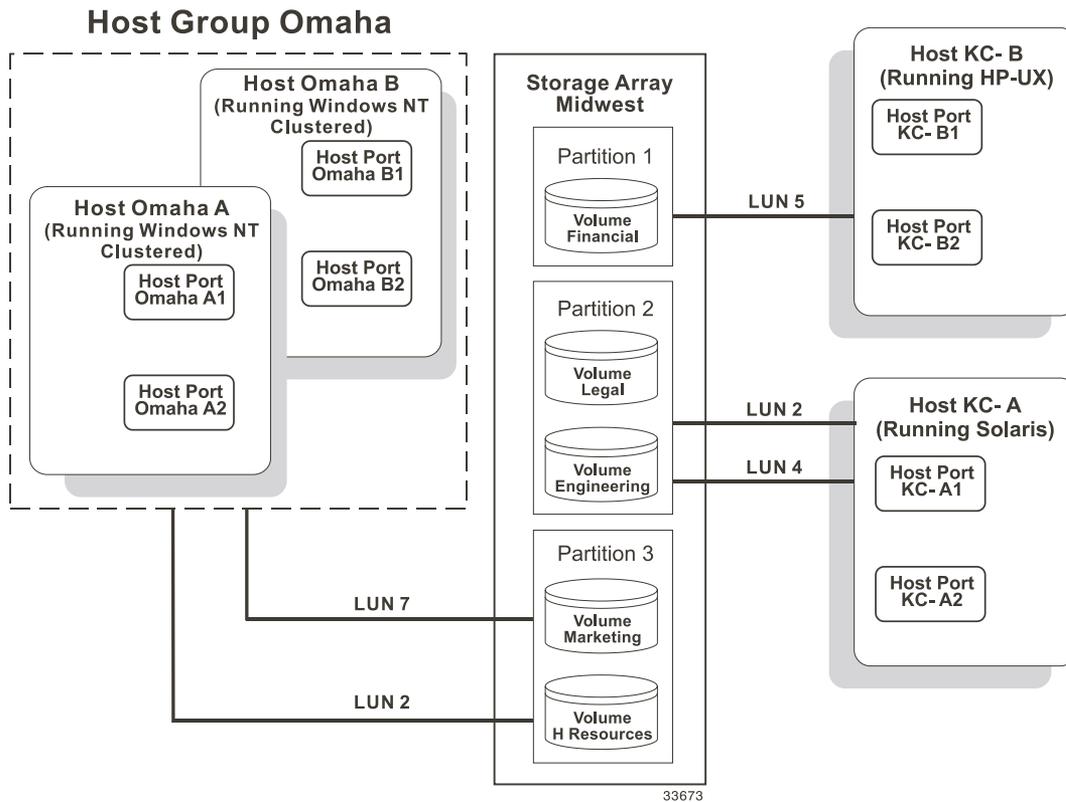


Figure 2-5 Heterogeneous Hosts Example

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 49 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 56).
- Monitor storage array performance and tune performance, if necessary, (as explained in “Tuning Storage Arrays” on page 61).

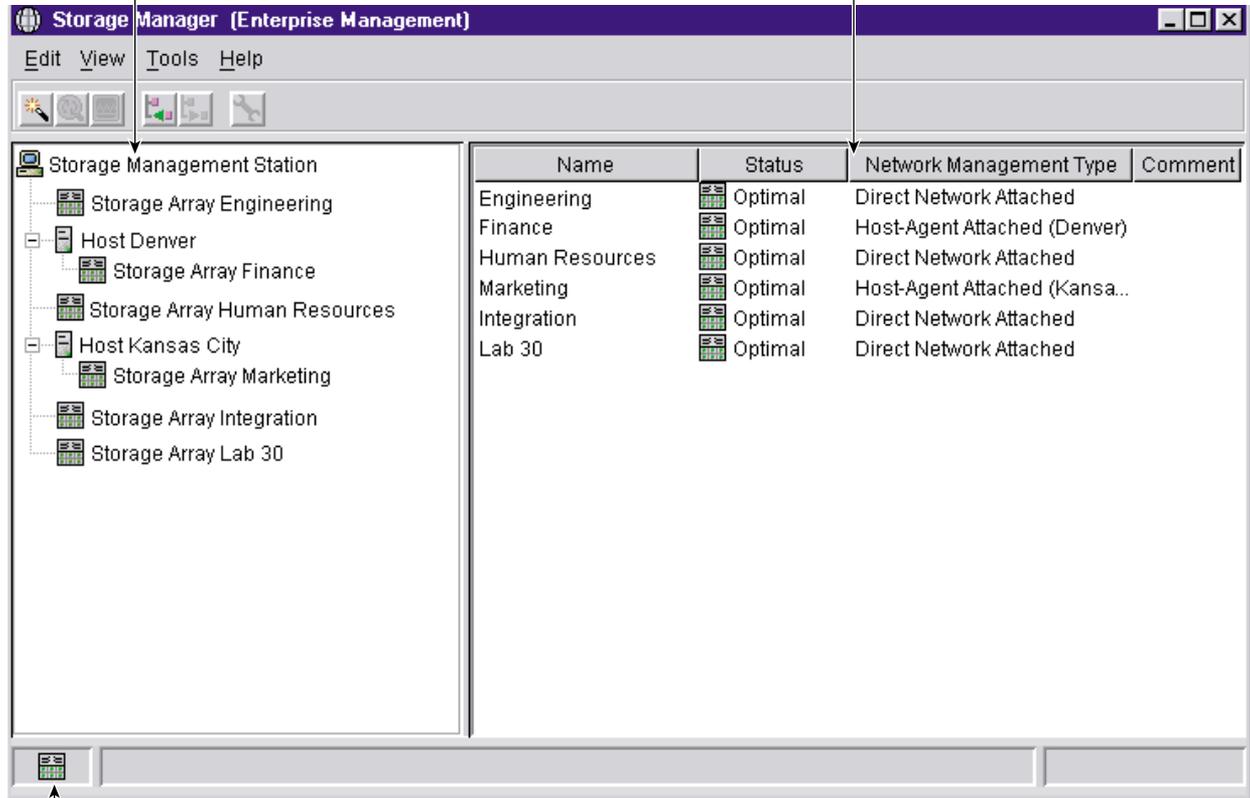
Monitoring the Health of Storage Arrays

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 2-5 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 2-5 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 2-5 Storage Array Status Indications (**continued**)

Status	Icon	Indication
Unresponsive		The storage management station cannot communicate with the only controller or both controllers over its network management connection to the storage array. 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.
Contacting Device		You have started the Enterprise Management window and the storage management software is establishing contact with the storage array. Note: This status is not displayed in the Logical view of the Array Management window.

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.

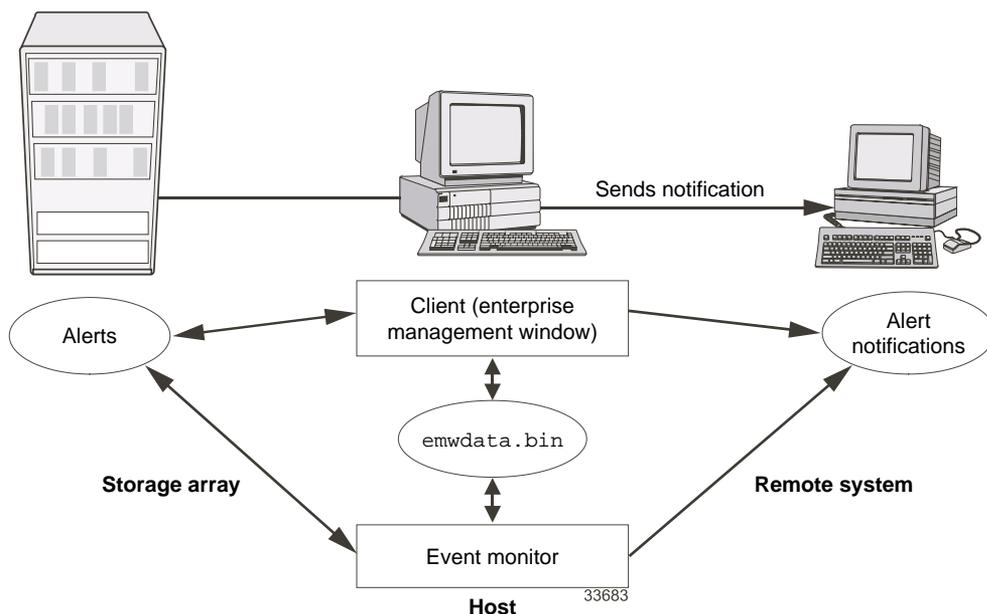


Figure 2-7 Event Monitor Configuration

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.

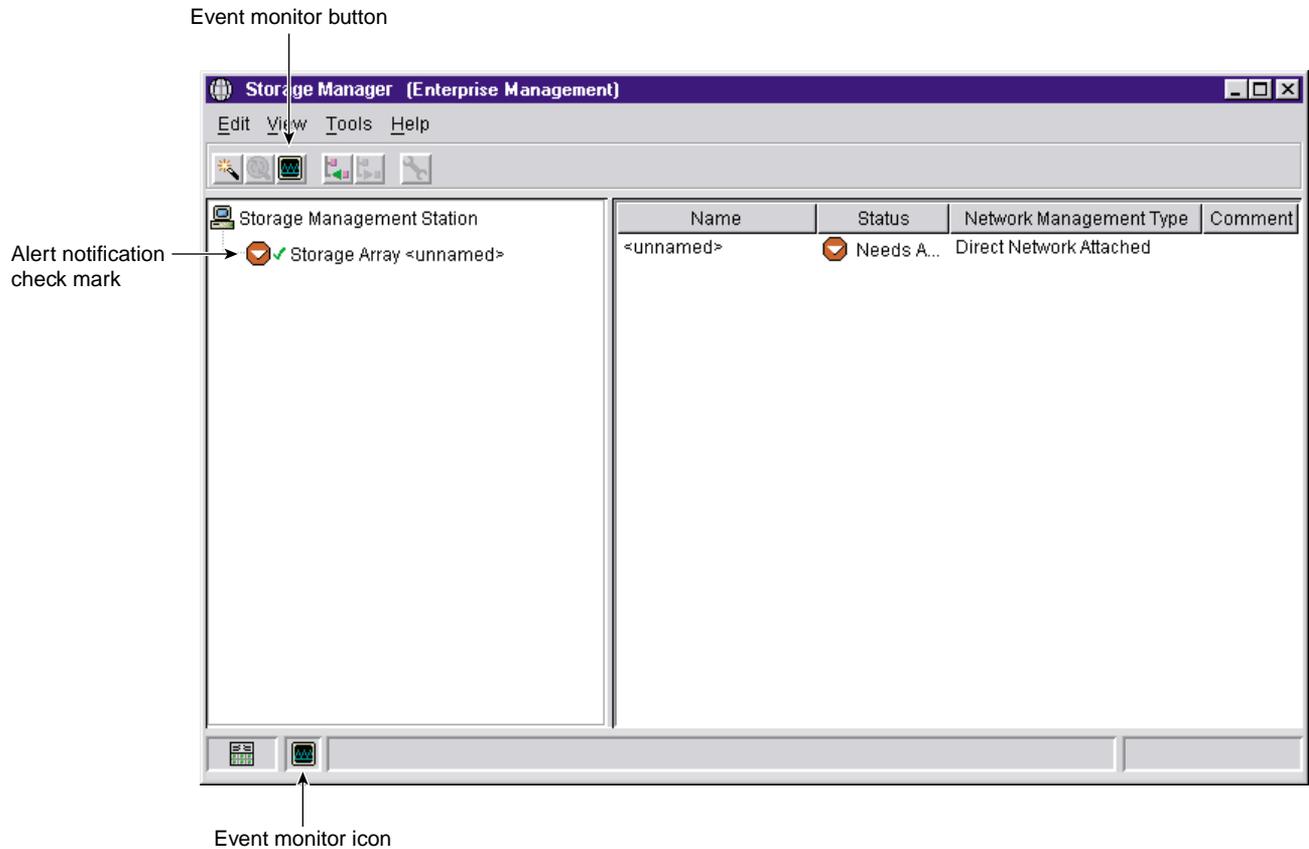


Figure 2-8 Event Monitor Example

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. See the *SGI TP9500 RAID Installation and Upgrade Guide* for more information.

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 `usr\`) 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**.

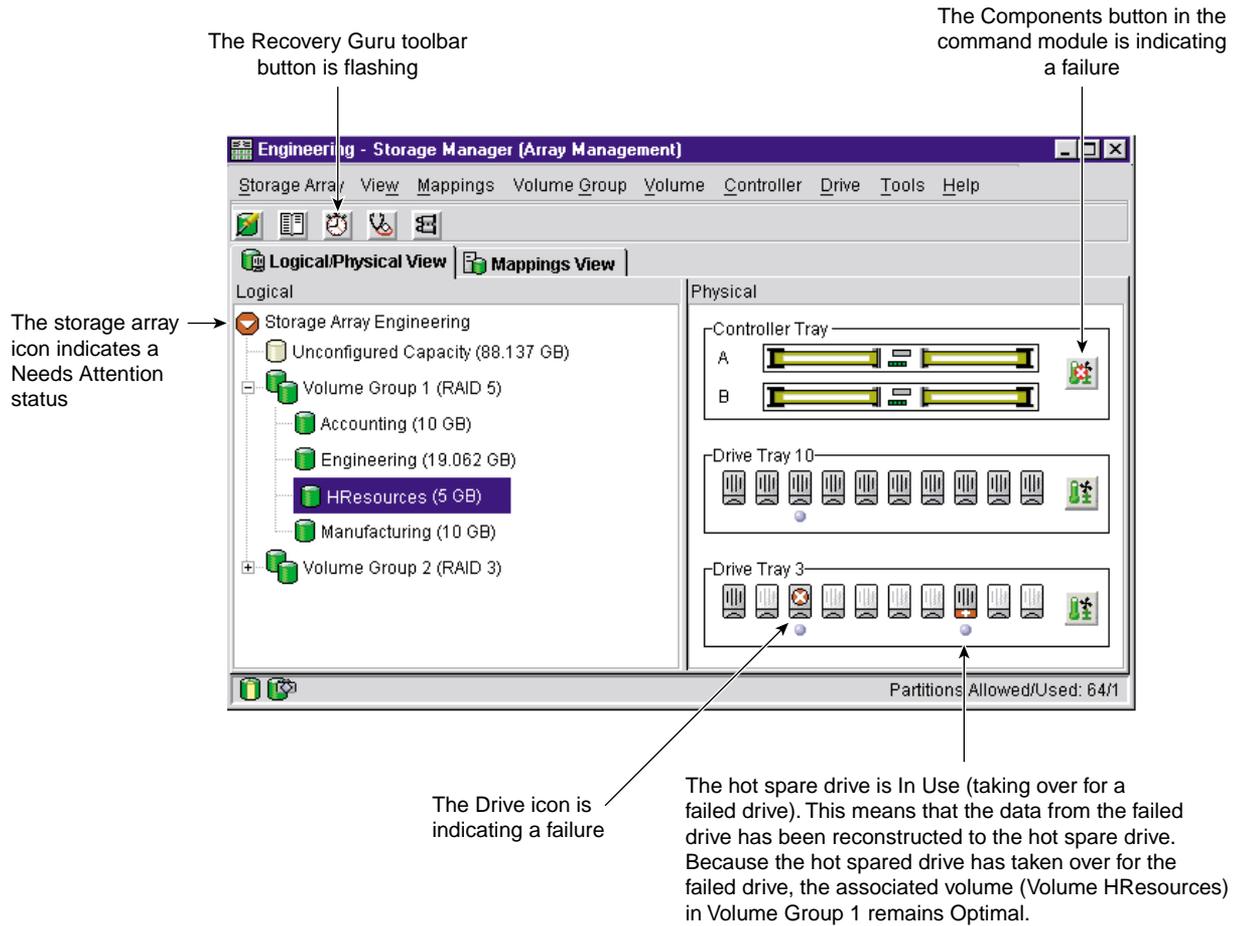


Figure 2-9 Problem Notification in the Array Management Window

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 TPSSM 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.

The Recovery Guru toolbar button

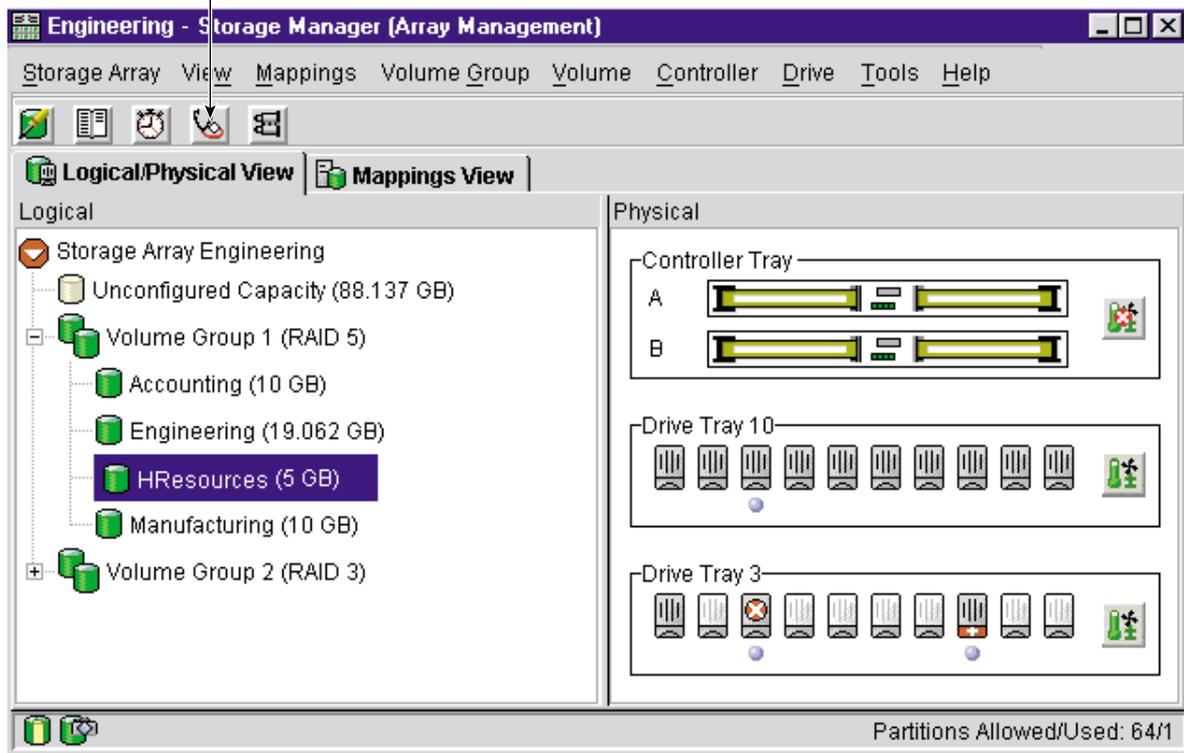


Figure 2-10 Recovery Guru Toolbar Button

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 informatin 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.

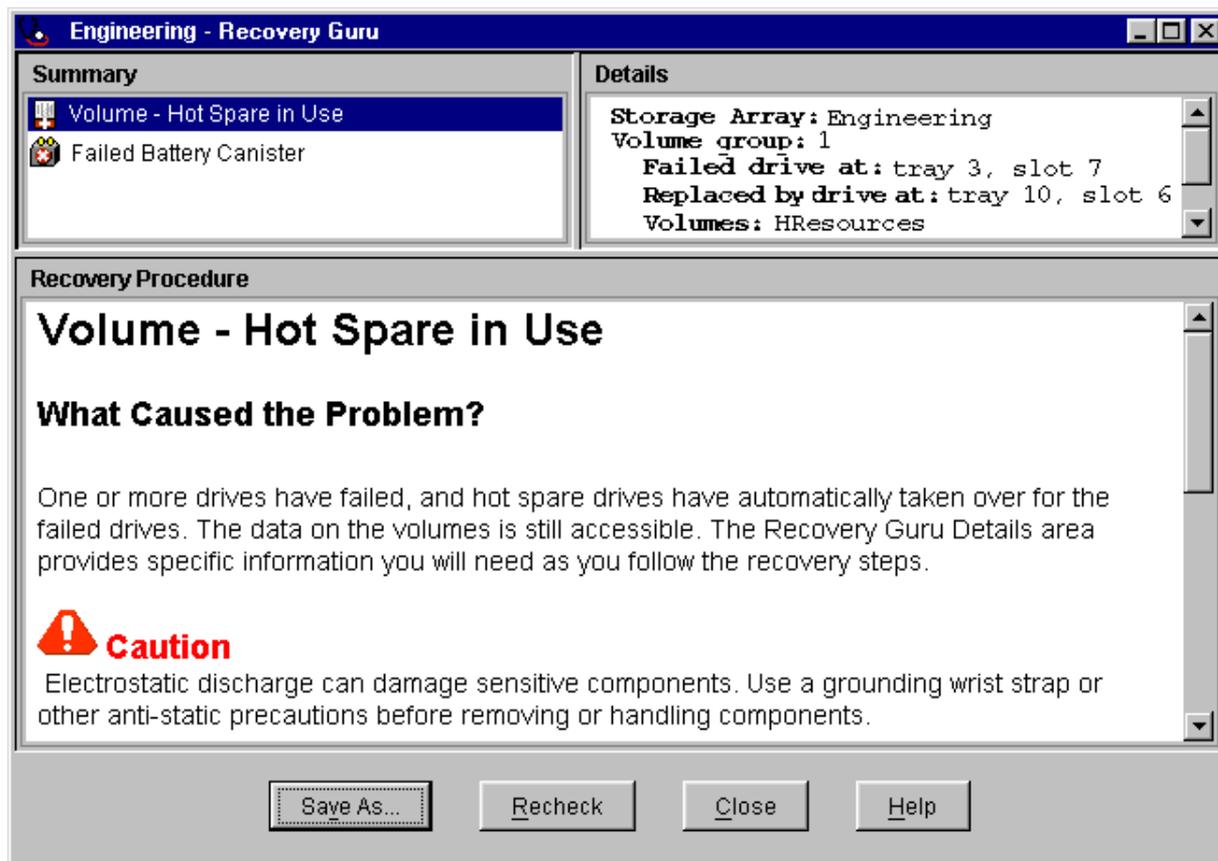


Figure 2-11 Recovery Guru Window

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.

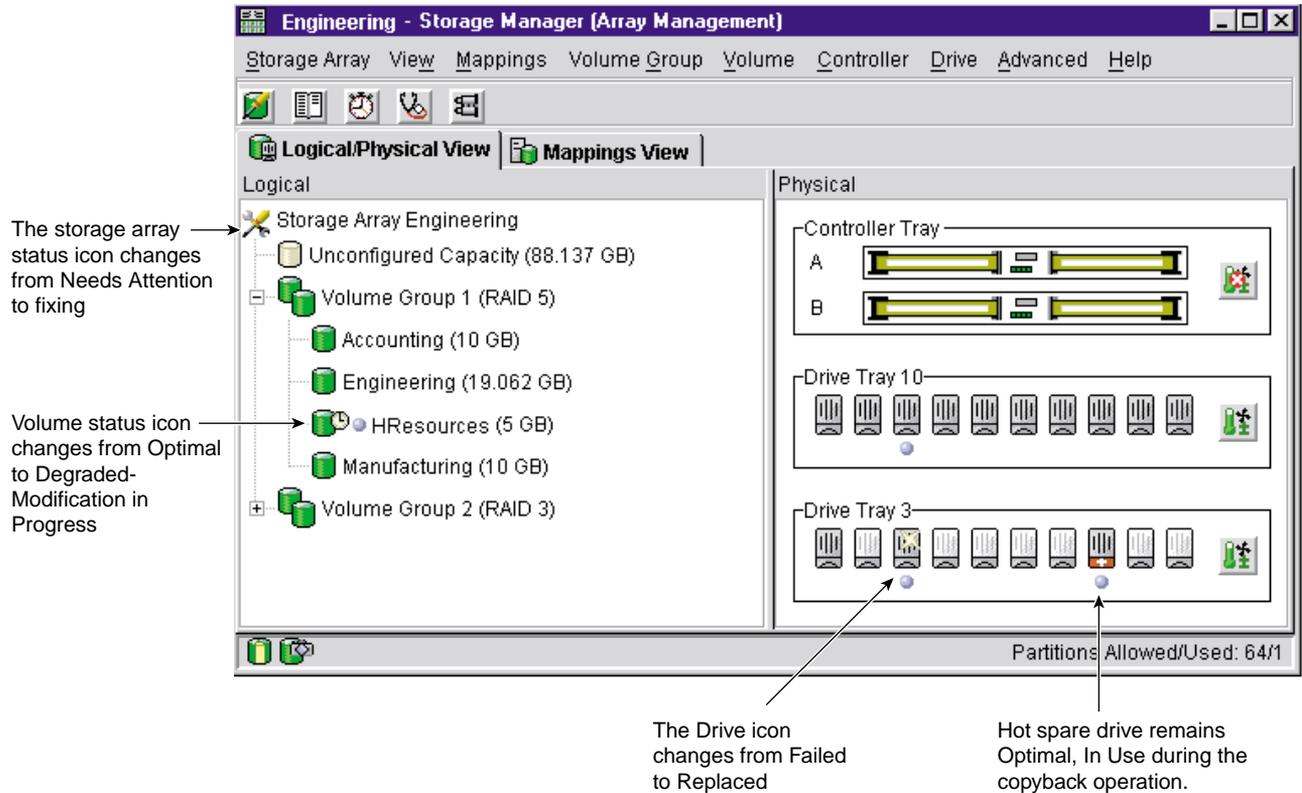


Figure 2-12 Status Changes During Example Recovery Operation

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

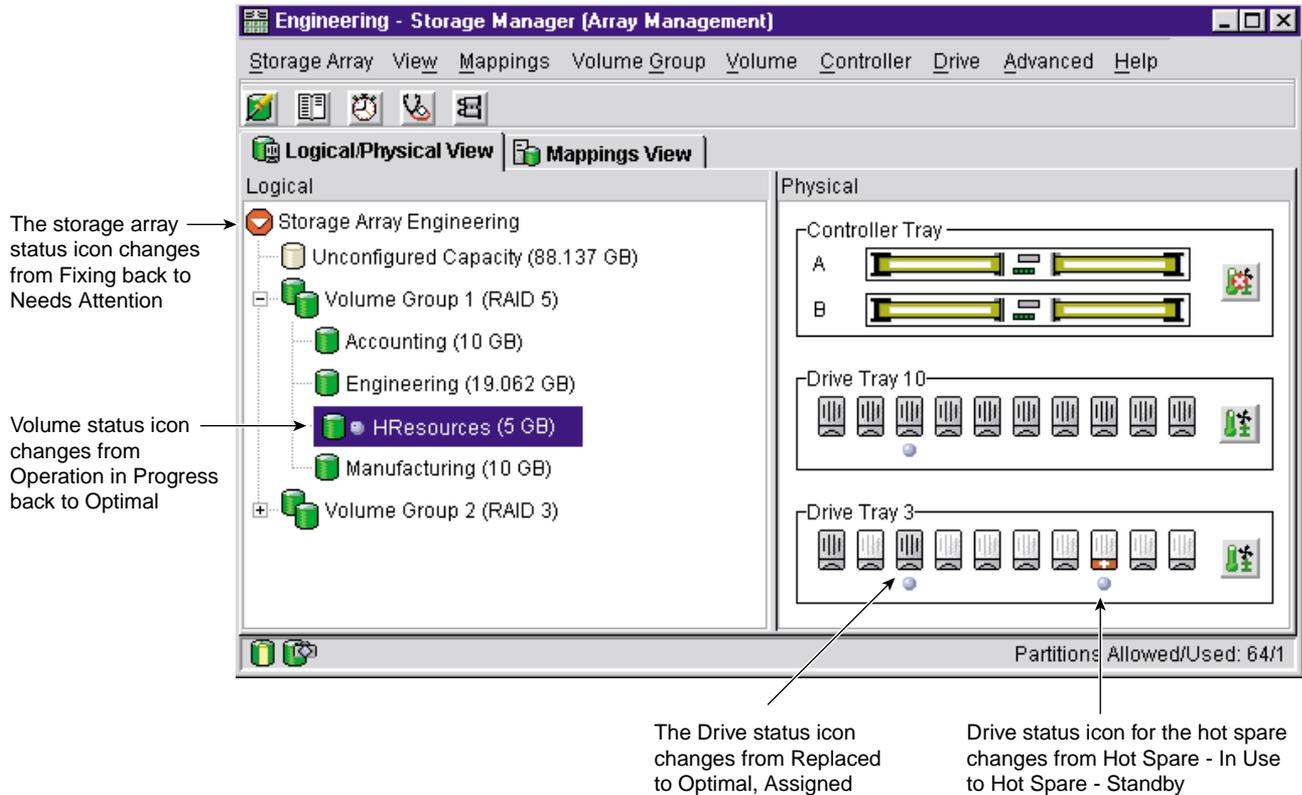


Figure 2-13 Status Changes When an Example Recovery Operation Completes

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 2-6.

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

Table 2-6 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 27 and the Array Management window Help.</p>

Table 2-6 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 2-6 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.</p> <p>Volume modification operations include reconstruction, copyback, initialization, defragmentation, change of RAID level, and change of segment size.</p>

Table 2-6 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 2-7 is a task-based index to the appropriate Help system.

Table 2-7 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 2-8 describes the Storage Array Configuration Help system.

Table 2-8 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 2-8 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 2-9 describes the Performance and Tuning Help system.

Table 2-9 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 2-10 describes the Event Notification Help system.

Table 2-10 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 2-11 describes the Recovering From Problems Help system.

Table 2-11 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 2-11 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 2-12 describes the Miscellaneous System Administration Help system.

Table 2-12 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 2-12 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 71
- “Hardware and Software Installation” on page 94
- “Using the Remote Volume Mirroring Premium Feature” on page 123
- “Disaster Recovery” on page 137
- “Troubleshooting” on page 145

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 71
- “Data Replication” on page 75
- “Connectivity and Input/Output” on page 78
- “Using Other Premium Features and Remote Volume Mirroring” on page 82
- “Hardware and Software Requirements” on page 86
- “Recommended Configurations” on page 88
- “Hardware and Software Checklist” on page 92

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.

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. There can be up to 8 defined mirrors on a SGI TP9400 storage array and up to 32 defined mirrors on a SGI TP9500 storage array.

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.

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.

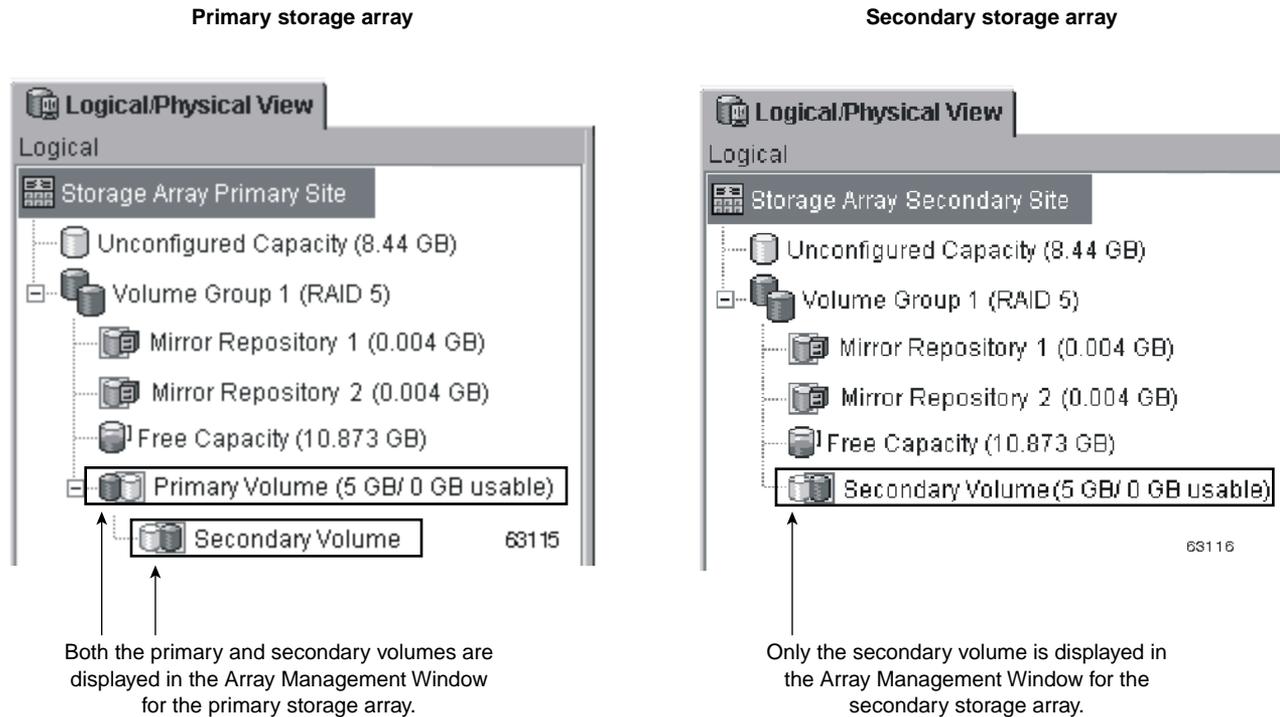


Figure 3-1 Primary and Secondary Volumes Displayed in the Array Management Window

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 4 MB, or 8 MB total for both mirror repository volumes of a dual controller array.

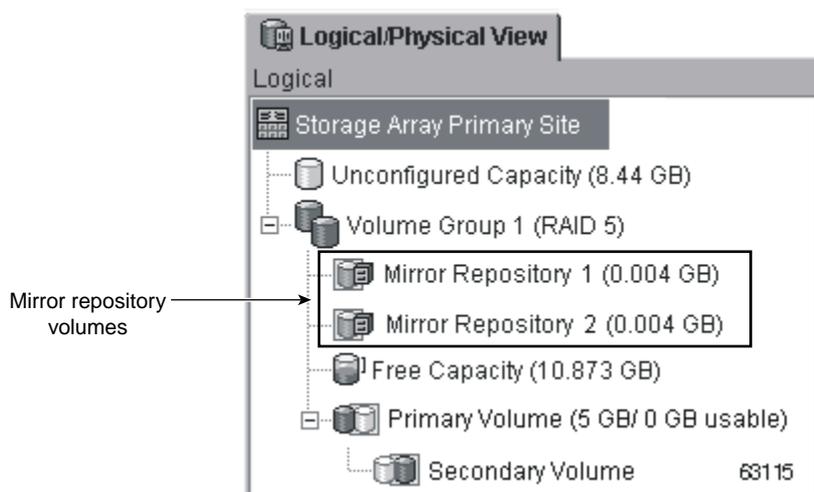


Figure 3-2 Mirror Repository Volumes Displayed in the **Array Management** Window

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 128.

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.

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 writes the data to the primary volume. The controller then initiates a remote write operation to copy the affected data blocks to the secondary volume at the secondary storage array.

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. See Figure 3-3 for a graphical representation of the data replication process.

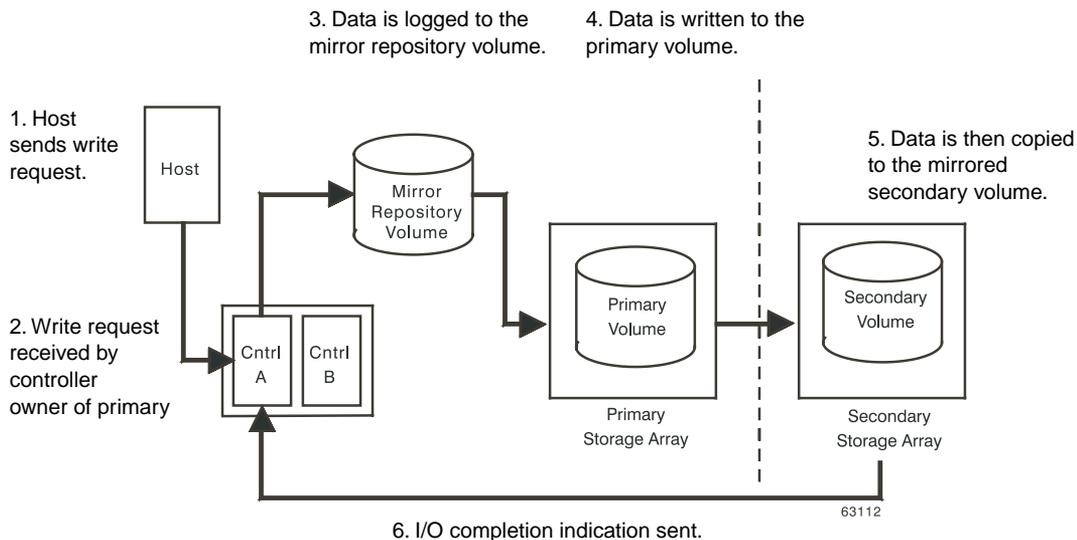


Figure 3-3 Data Replication between Storage Arrays

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.

Link Interruptions or Secondary Volume Errors

When processing write requests, the primary controller may 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 and secondary volumes are no longer correctly mirrored. The primary controller transitions the mirrored pair into an Unsynchronized status and sends an I/O completion to the primary host, as shown in Figure 3-4. The primary host can continue to write to the primary volume but remote writes will not take place.

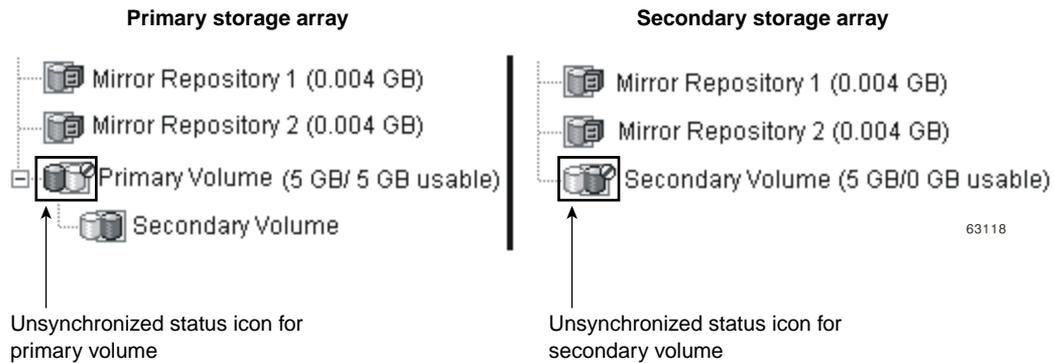


Figure 3-4 Mirrored Volume Pair in Unsyncronized Status

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 Unsyncronized status to a Synchronization-in-Progress status, as shown in Figure 3-5.

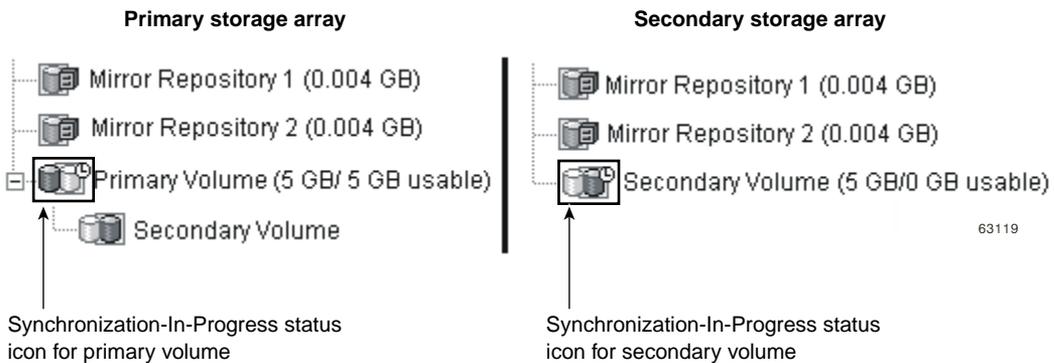


Figure 3-5 Mirrored Volume Pair in Synchronization-in-Progress Status

The primary controller will also mark the mirrored pair as Unsyncronized 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 Unsyncronized. 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 123.

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.

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 88.
- 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 90.

- 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 91.

For instructions on how to configure storage arrays, see “Hardware and Software Installation” on page 94.

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.

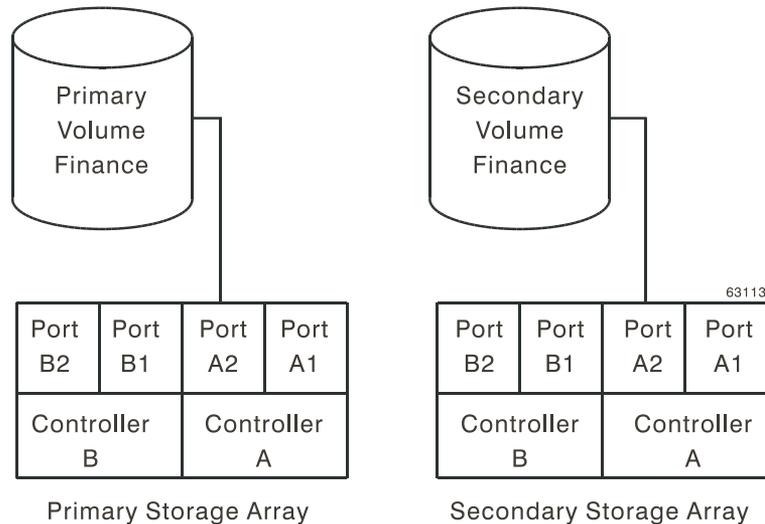


Figure 3-6 Controller A as Controller Owner

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.

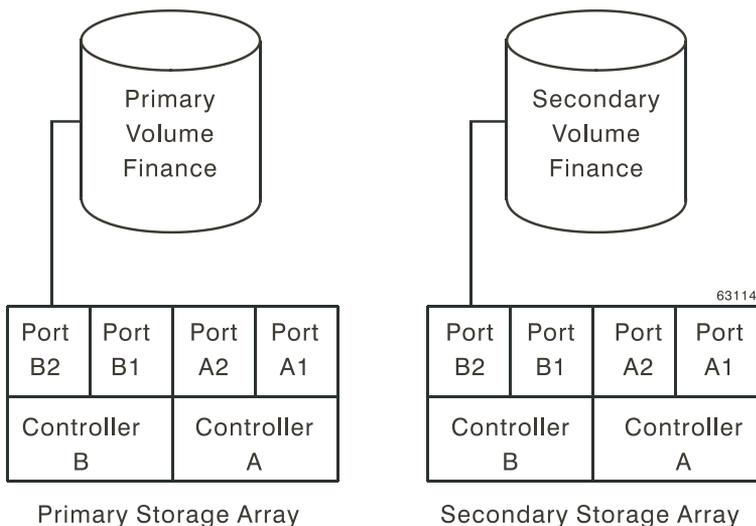


Figure 3-7 Controller B as Controller Owner

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 between storage arrays can be achieved.

Note: Specialized third-party products allow for greater distances to be supported between storage arrays in a Fibre Channel environment. However, in this version of the storage management software, distances between storage arrays have been tested and verified up to 10 km.

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 123.

- 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.

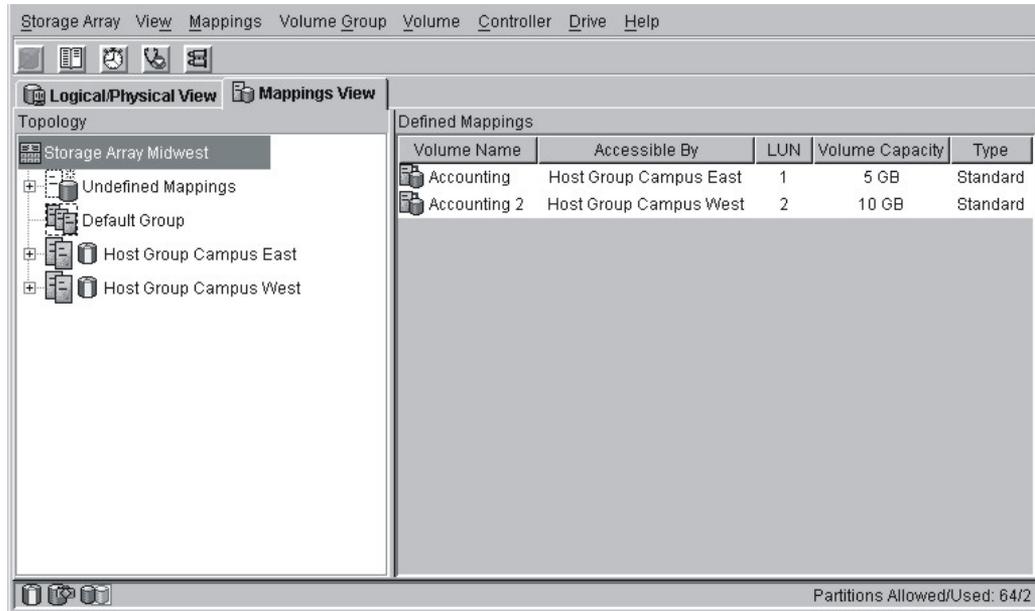
Using Other Premium Features and Remote Volume Mirroring

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.

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 a partition accessible by Host Group Campus West.



63117

Figure 3-8 Storage Partitioning Example

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.

Snapcopy Volumes

A snapcopy 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 snapcopy 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 snapcopy volume is permitted to be a candidate for the primary volume in a mirror relationship. The snapcopy repository volume stores

information about all data altered since the snapcopy was created. Figure 3-9 shows the primary volume (also the base volume for the snapcopy volume), the secondary volume, and the snapcopy volume in the **Array Management** window for the primary storage array.

Note: The secondary volume of a mirror relationship **cannot** be a base volume with an associated snapcopy volume.

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

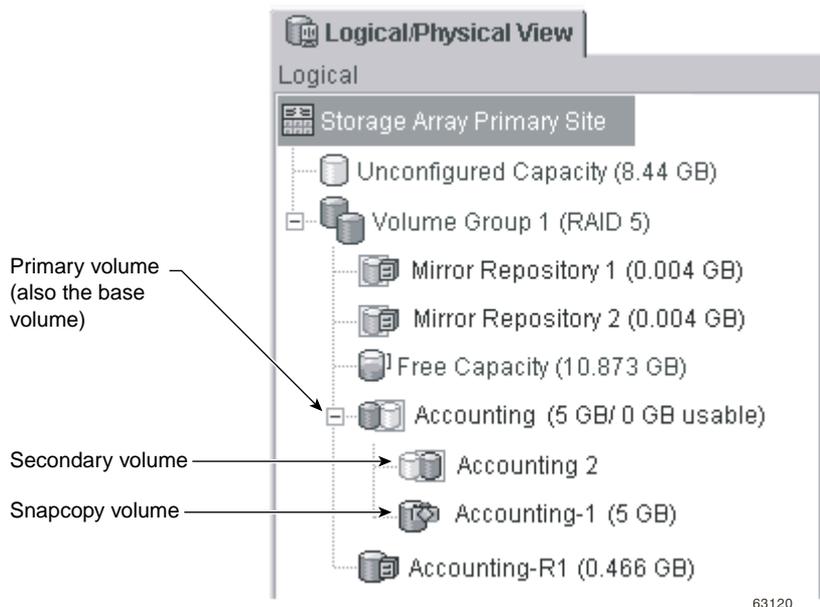


Figure 3-9 Snapcopy Volume Included in a Mirror Relationship

Snapcopy Volume Failure

In this version of the storage management software, snapcopy 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 snapcopy volumes associated with the demoted primary volume will be failed.

Note: A secondary volume **cannot** be selected to be the base volume of a snapcopy 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 snapcopy volume of Primary Volume Accounting has been created.

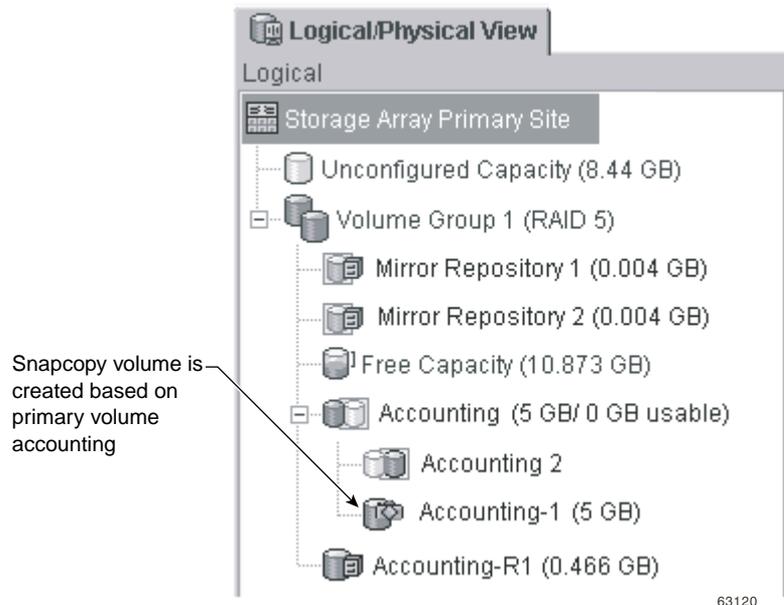


Figure 3-10 Snapcopy Volume Associated with Primary Volume Accounting

If a role reversal is performed, Primary Volume Accounting is demoted to a secondary role and the snapcopy 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 snapcopy volume, which is not permitted.

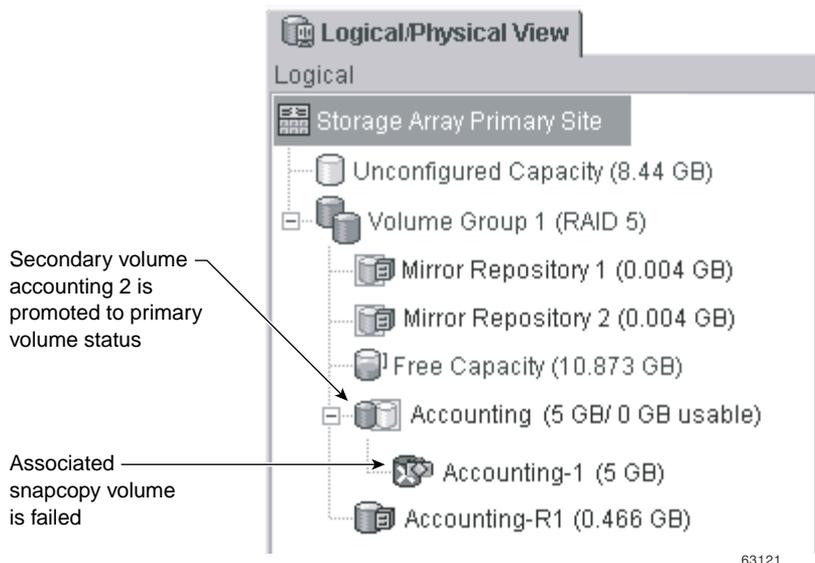


Figure 3-11 Snapcopy Volume Failure Because of Role Reversal

If a snapcopy 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 snapcopy volume will return to an Optimal status. For more information, see “Troubleshooting” on page 145.

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 supplementary to the standard required for a basic storage array environment as described in the *SGI TP9500 RAID Installation and Upgrade Guide*.

- 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.

The certified switches for use with the Remote Volume Mirroring feature are as follows.

- **Brocade Silkworm 2400** – 8-port, 1 Gigabit/sec (Highest Availability Campus Configuration and Intra-Site Configuration only)
- **Brocade Silkworm 2800** – 16-port, 1 Gigabit/sec
- **Brocade Silkworm 3800** – 16-port, 2 Gigabit/sec
- **Brocade Silkworm 3200** – 8-port, 2 Gigabit/sec
- 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 8.30 or later
- Controller Firmware Version System Manager
- Remote Volume Mirroring Feature Key

Certified Host Operating Systems

Hosts accessing mirrored volumes should be running one of the following certified operating systems when using the Remote Volume Mirroring feature:

- Windows NT 4.0 SP 6a Server Edition
- Windows 2000 SP 2
- Windows 98
- Solaris 2.6, 2.7 (Solaris 7), 2.8 (Solaris 8)
- HP-UX 11.0 or later
- Linux (Red Hat 6.x)
- Netware 5.1
- IRIX 6.5.14 or later
- AIX 4.3.3 or later

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

The Highest Availability Campus Configuration is a fully redundant configuration, and consists of two storage arrays and four switches connected using a Fibre Channel fabric as shown in Figure 3-12. 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 94.

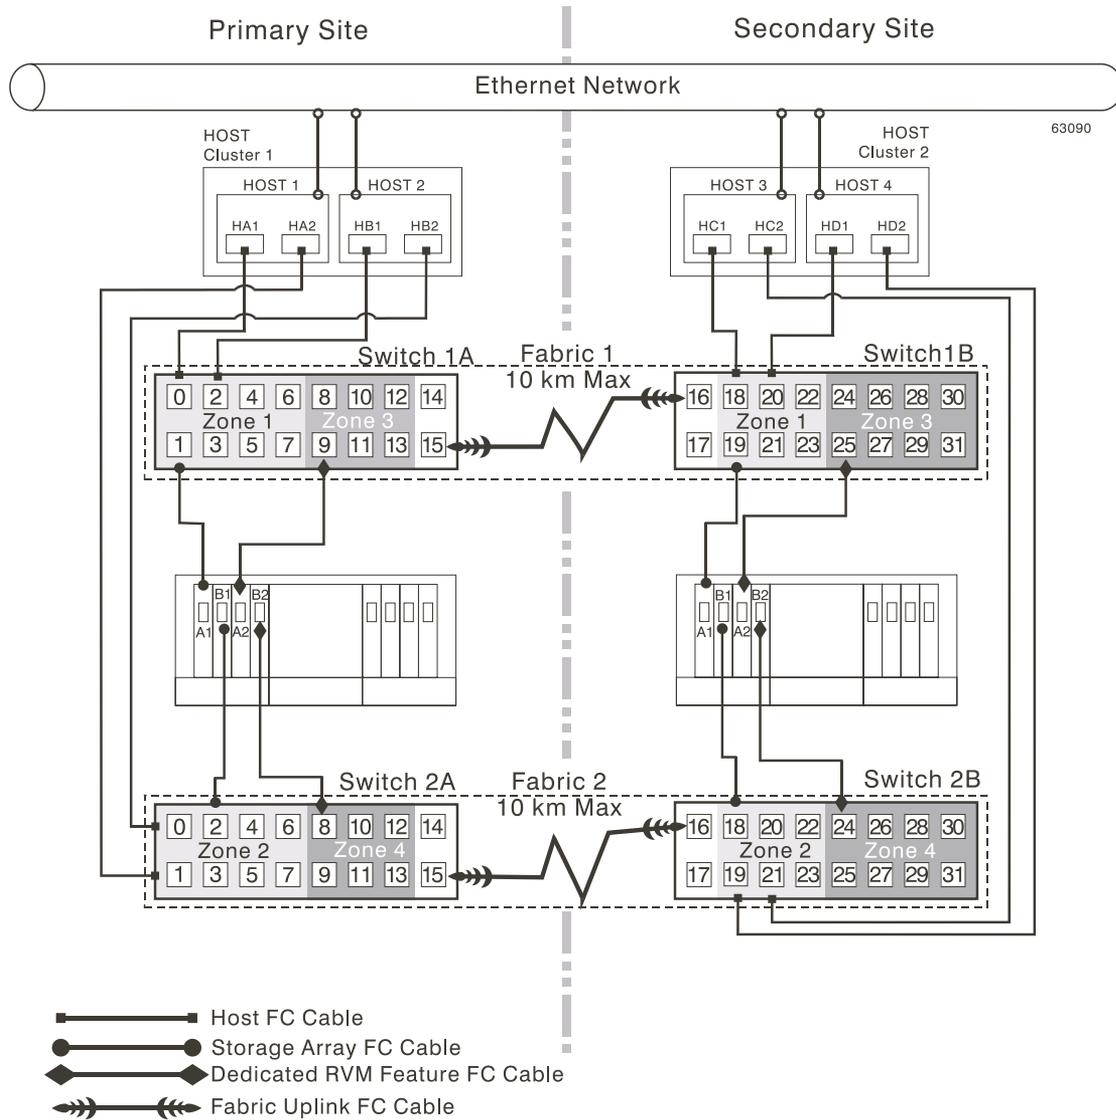


Figure 3-12 Highest Availability Campus Configuration

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 94.

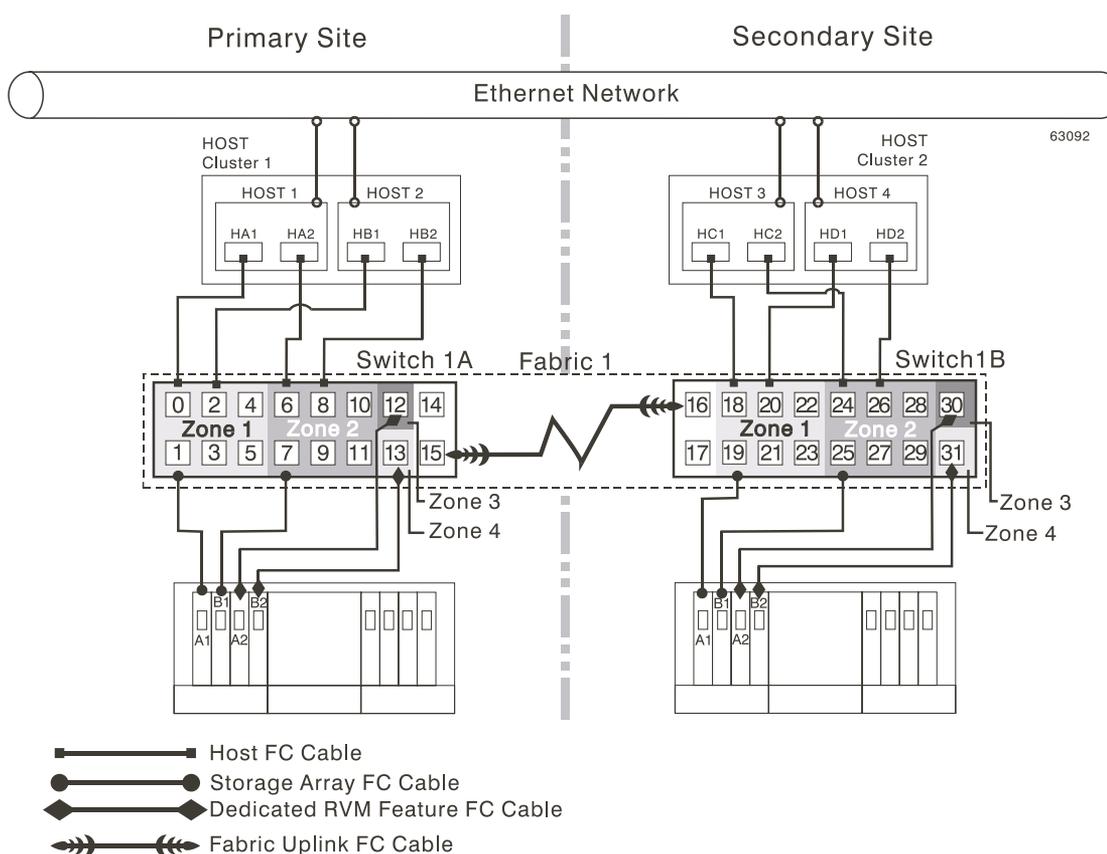


Figure 3-13 Campus Configuration

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 94.

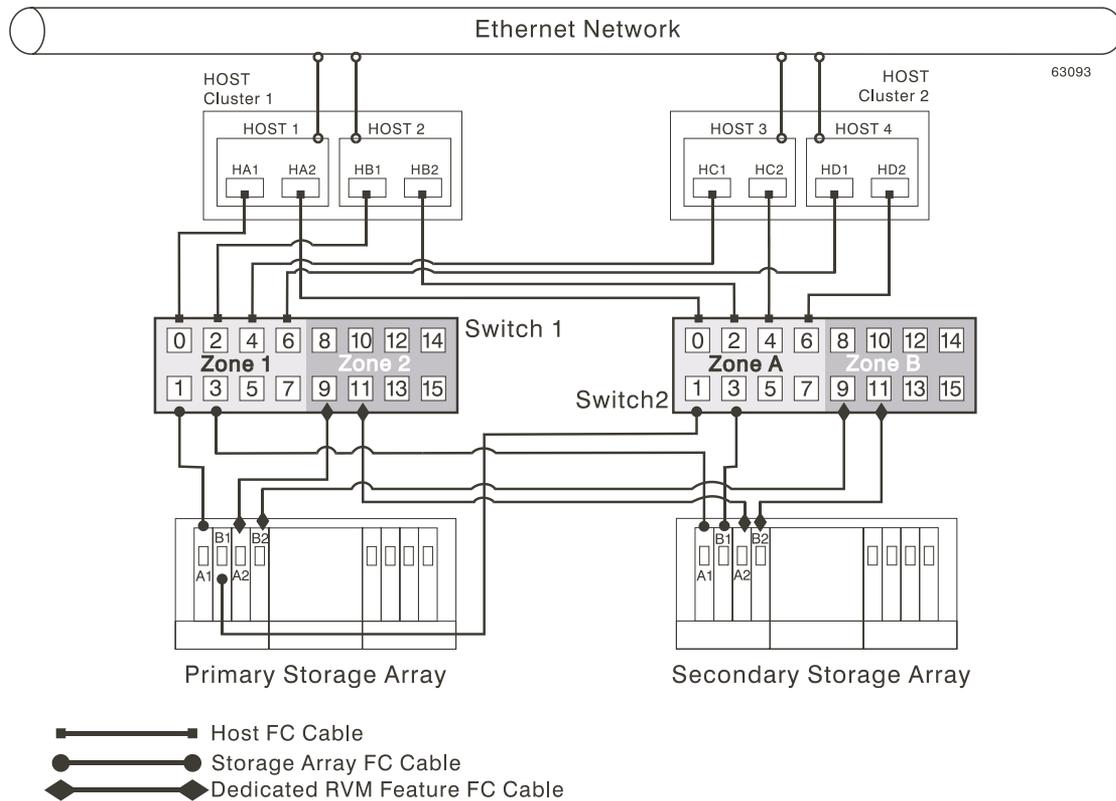


Figure 3-14 Intra-Site Configuration

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.	<input type="checkbox"/> Completed.
For a complete list of the hardware requirements, see "Hardware Requirements" on page 87.	
2. Ensure your primary and secondary storage arrays have been properly configured.	<input type="checkbox"/> Completed.
See the <i>SGI TP9500 RAID Installation and Upgrade Guide</i> for storage array configuration information.	
3. Ensure that your switches and cables have been properly configured.	<input type="checkbox"/> Completed.
For more information on configuring switches for use with the Remote Volume Mirroring feature, see "Hardware and Software Installation" on page 94.	

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 94.	
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 94 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 123.	
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 123.	

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 94
- “Hardware Installation” on page 97
- “Software Installation” on page 121

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 basic site preparation considerations including the use of Fibre Channel cables and connections, see the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.

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 94.

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: The uplink port that connects (cascades) switches within a fabric should not be zoned.

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.

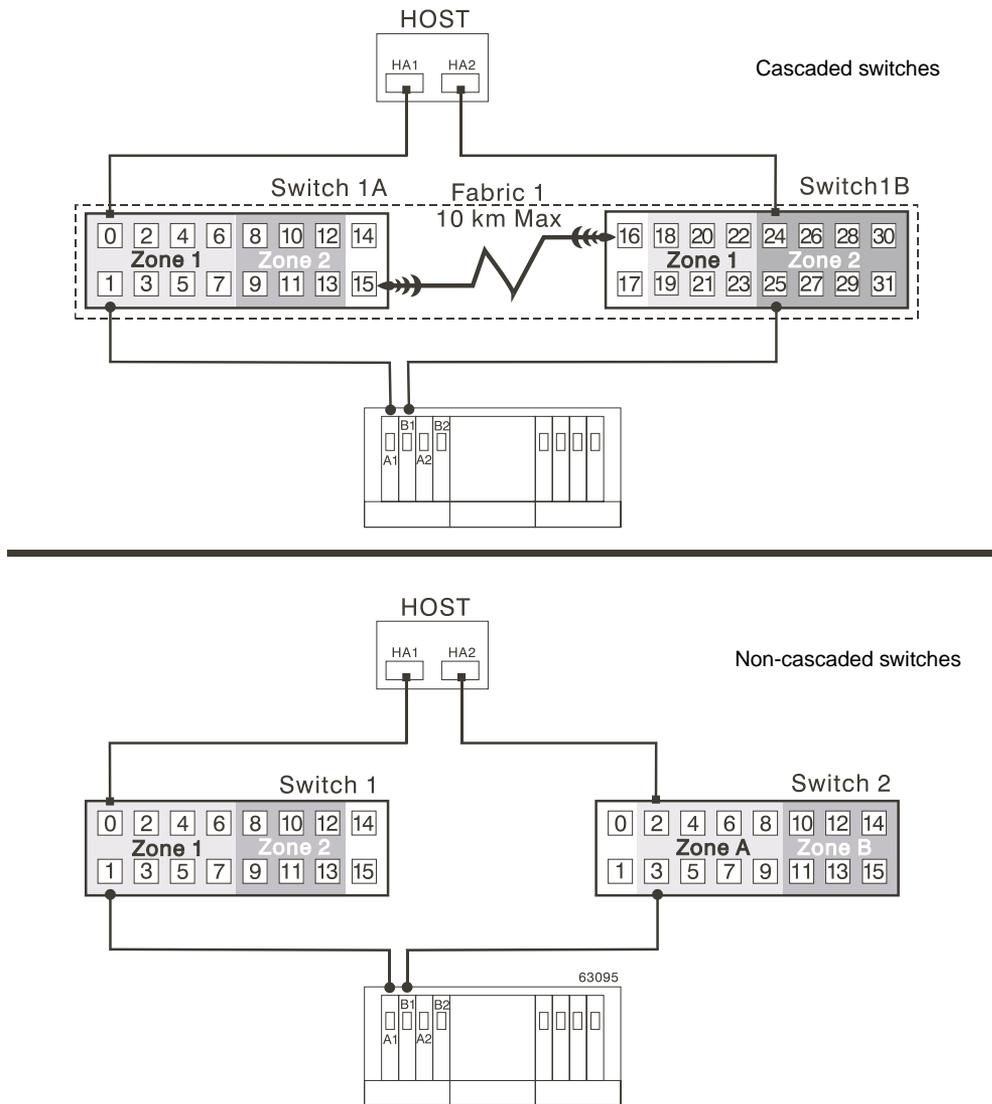


Figure 3-15 Switch Zoning in Cascaded and Non-Cascaded Fibre Channel Switches

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 proceeding each configuration 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 88.
- **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 90.
- **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 91.

For a list of required materials to complete any hardware installation, see “Hardware and Software Installation” on page 94.

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.

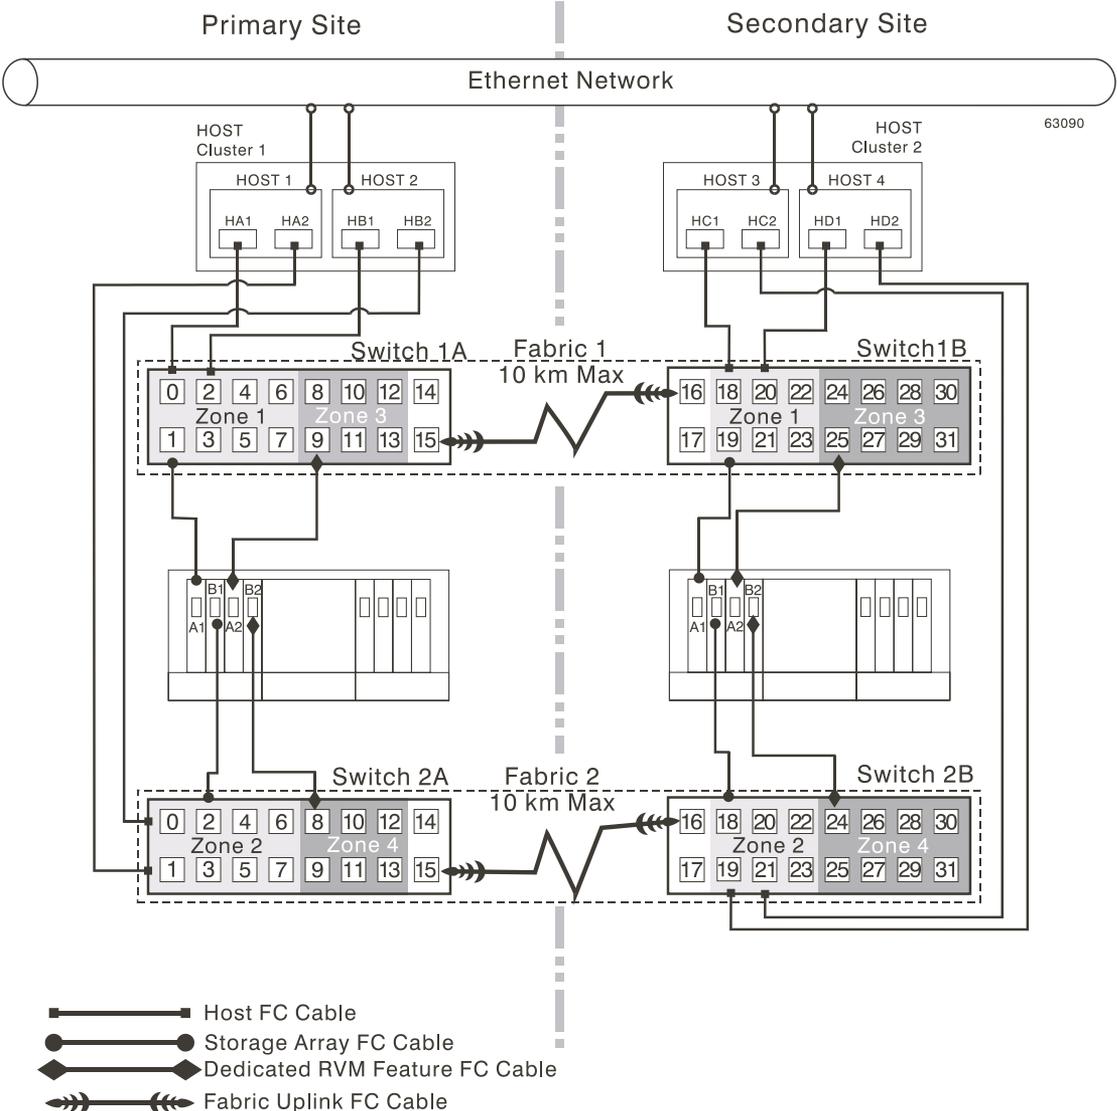


Figure 3-16 Highest Availability Campus Configuration

Switch 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 will be 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.
- The uplink port must not be zoned on any Fibre Channel switch.
- Switches are zoned so that a single host adapter can access only one controller per storage array.

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.

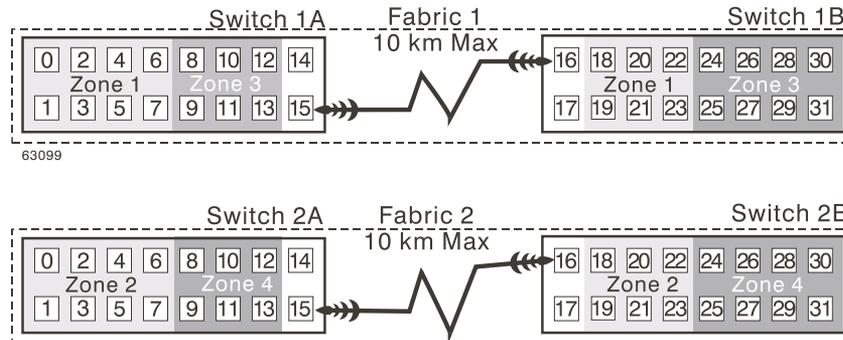


Figure 3-17 Switch Zoning for the Highest Availability Campus Configuration

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 95.

Cabling for the Highest Availability Campus Configuration

Table 3-3 is provided for a convenient reference of port connections.

Table 3-3 Highest Availability Campus Configuration Connections

Site	Switch	Zone	Port	Connect To/From
Primary	1A	1	0	Host 1, HBA 1 (HA1)
Primary	1A	1	1	Controller Port A1
Primary	1A	1	2	Additional Host, HBA 1 (Optional)
Primary	1A	1	3	Additional Host, HBA 1 (Optional)
Primary	1A	1	4	Additional Host, HBA 1 (Optional)
Primary	1A	1	5	Additional Host, HBA 1 (Optional)
Primary	1A	1	6	Additional Host, HBA 1 (Optional)
Primary	1A	1	7	Additional Host, HBA 1 (Optional)
Primary	1A	3	8	Unused

Table 3-3 Highest Availability Campus Configuration Connections (**continued**)

Site	Switch	Zone	Port	Connect To/From
Primary	1A	3	9	Controller Port A2 (Dedicated RVM Port)
Primary	1A	3	10	Unused
Primary	1A	3	11	Unused
Primary	1A	3	12	Unused
Primary	1A	3	13	Unused
Primary	1A	None	14	Spare (May be added to either Zone)
Primary	1A	None	15	Uplink with Switch 1B
Primary	2A	2	0	Host 2, HBA 2 (HB2)
Primary	2A	2	1	Host 1, HBA 2 (HB2)
Primary	2A	2	2	Controller Port B1
Primary	2A	2	3	Additional Host, HBA 2 (Optional)
Primary	2A	2	4	Additional Host, HBA 2 (Optional)
Primary	2A	2	5	Additional Host, HBA 2 (Optional)
Primary	2A	2	6	Additional Host, HBA 2 (Optional)
Primary	2A	2	7	Additional Host, HBA 2 (Optional)
Primary	2A	4	8	Controller Port B2 (Dedicated RVM Port)
Primary	2A	4	9	Unused
Primary	2A	4	10	Unused
Primary	2A	4	11	Unused
Primary	2A	4	12	Unused
Primary	2A	4	13	Unused
Primary	2A	None	14	Spare (May be added to either Zone)
Primary	2A	None	15	Uplink with Switch 2B
Secondary	1B	None	16	Uplink with Switch 1A

Table 3-3 Highest Availability Campus Configuration Connections **(continued)**

Site	Switch	Zone	Port	Connect To/From
Secondary	1B	None	17	Spare (May be added to either Zone)
Secondary	1B	1	18	Host 3, HBA 1 (HC1)
Secondary	1B	1	19	Controller Port A1
Secondary	1B	1	20	Host 4, HBA 1 (HD1)
Secondary	1B	1	21	Additional Host, HBA 1 (Optional)
Secondary	1B	1	22	Additional Host, HBA 1 (Optional)
Secondary	1B	1	23	Additional Host, HBA 1 (Optional)
Secondary	1B	3	24	Unused
Secondary	1B	3	25	Controller Port A2 (Dedicated RVM Port)
Secondary	1B	3	26	Unused
Secondary	1B	3	27	Unused
Secondary	1B	3	28	Unused
Secondary	1B	3	29	Unused
Secondary	1B	3	30	Unused
Secondary	1B	3	31	Unused
Secondary	2B	None	16	Uplink with Switch 2A
Secondary	2B	None	17	Spare (May be added to either Zone)
Secondary	2B	2	18	Controller Port B1
Secondary	2B	2	19	Host 4, HBA 2 (HD2)
Secondary	2B	2	20	Additional Host, HBA 2 (Optional)
Secondary	2B	2	21	Host 3, HBA 2 (HC2)
Secondary	2B	2	22	Additional Host, HBA 2 (Optional)
Secondary	2B	2	23	Additional Host, HBA 2 (Optional)
Secondary	2B	4	24	Controller Port B2 (Dedicated RVM Port)

Table 3-3 Highest Availability Campus Configuration Connections (**continued**)

Site	Switch	Zone	Port	Connect To/From
Secondary	2B	4	25	Unused
Secondary	2B	4	26	Unused
Secondary	2B	4	27	Unused
Secondary	2B	4	28	Unused
Secondary	2B	4	29	Unused
Secondary	2B	4	30	Unused
Secondary	2B	4	31	Unused

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.
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 as presented in the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.

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).

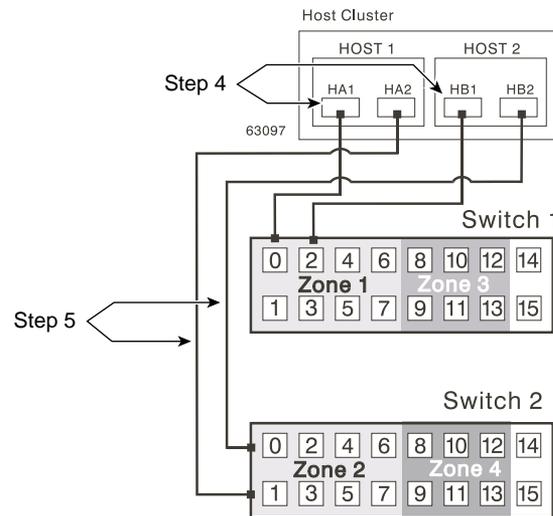


Figure 3-18 Host Bus Adapter Connections to Fibre Channel Switches

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 78.

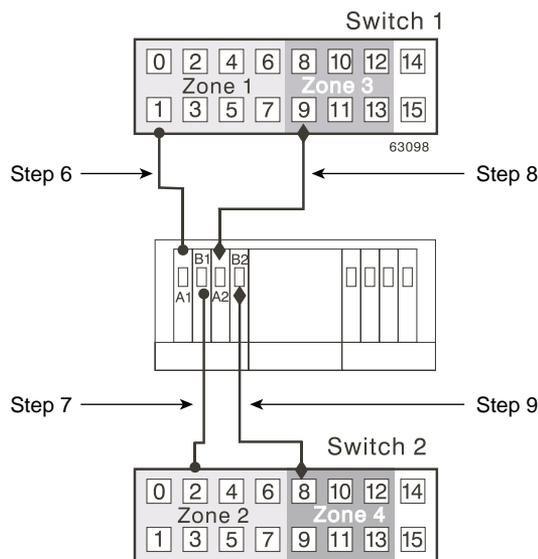


Figure 3-19 Storage Array Connections to Fibre Channel Switches

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 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).

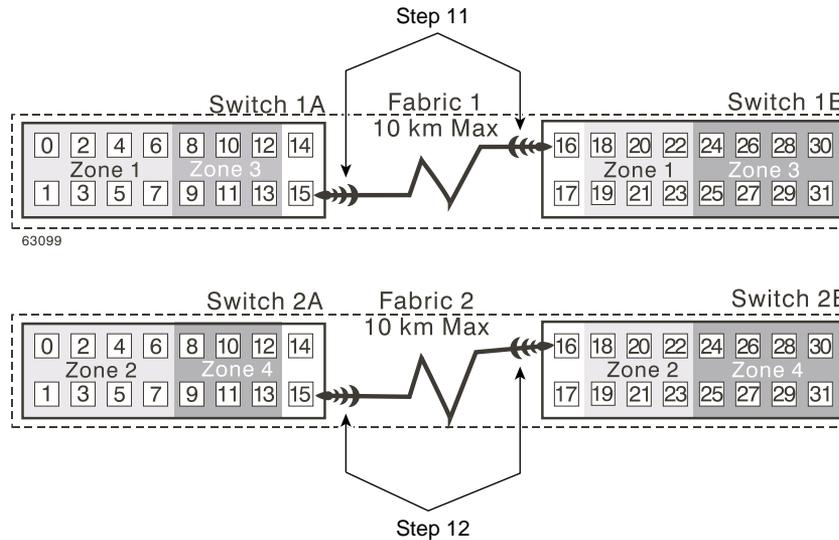


Figure 3-20 Connecting Switches over a Remote Network to Complete Fabric Environments

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.
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 121.

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.

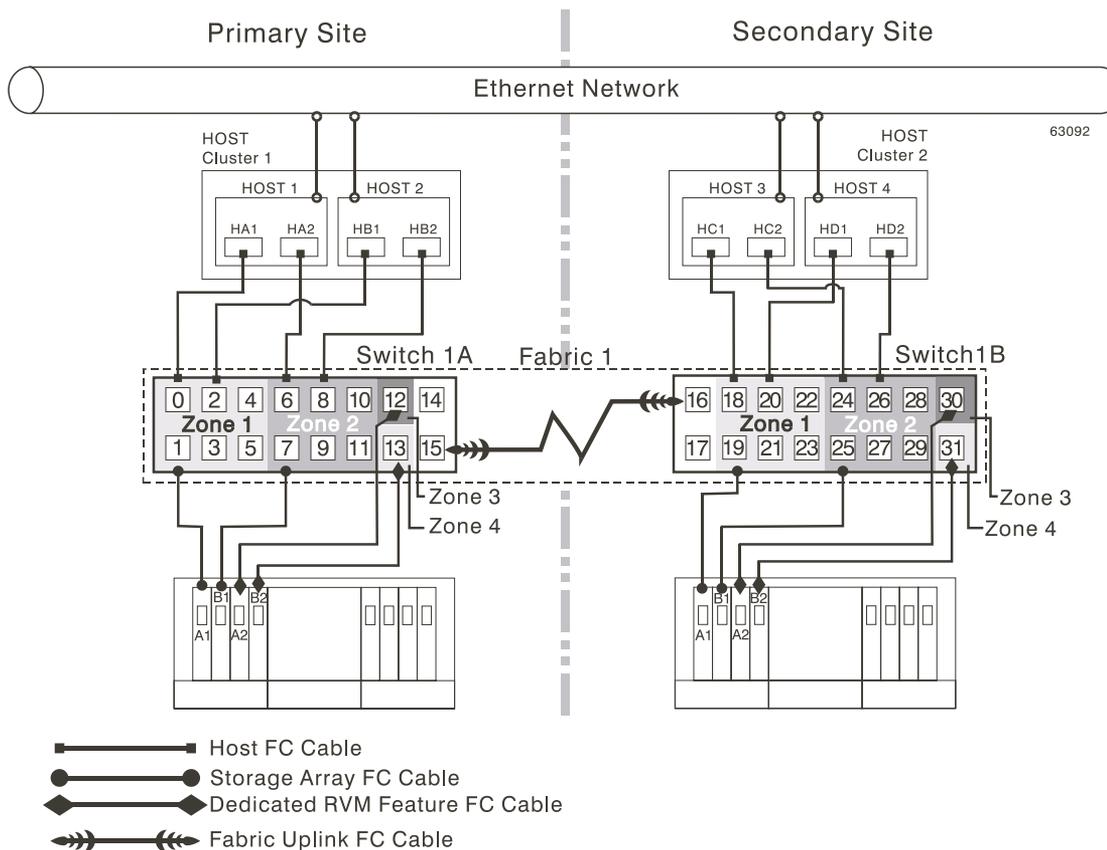


Figure 3-21 Campus Configurations

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.
- The uplink port must not be zoned on either Fibre Channel switch.
- 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.

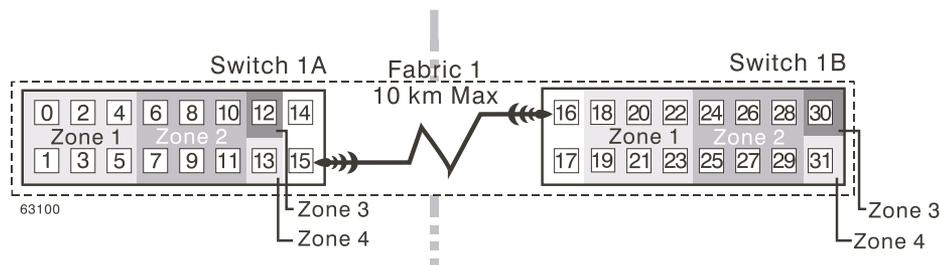


Figure 3-22 Switch Zoning for the Campus Configuration

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 95.

Cabling for the Campus Configuration

Table 3-4 is provided for a convenient reference of port connections..

Table 3-4 Campus Configuration Connections

Site	Switch	Zone	Port	Connect To/From
Primary	1A	1	0	Host 1, HBA 1 (HA1)
Primary	1A	1	1	Controller Port A1
Primary	1A	1	2	Host 2, HBA 1 (HB1)
Primary	1A	1	3	Additional Host, HBA 1 (Optional)
Primary	1A	1	4	Additional Host, HBA 1 (Optional)
Primary	1A	1	5	Additional Host, HBA 1 (Optional)
Primary	1A	2	6	Host 1, HBA 2 (HA2)
Primary	1A	2	7	Controller Port B1
Primary	1A	2	8	Host 2, HBA 2 (HB2)
Primary	1A	2	9	Additional Host, HBA 2 (Optional)
Primary	1A	2	10	Additional Host, HBA 2 (Optional)
Primary	1A	2	11	Additional Host, HBA 2 (Optional)
Primary	1A	3	12	Controller Port B2 (Dedicated RVM Port)
Primary	1A	4	13	Controller Port B2 (Dedicated RVM Port)
Primary	1A	None	14	Spare (May be added to any Zone)
Primary	1A	None	15	Uplink with Switch 1B
Secondary	1B	None	16	Uplink with Switch 1A
Secondary	1B	None	17	Spare (May be added to any Zone)
Secondary	1B	1	18	Host 3, HBA 1 (HC1)
Secondary	1B	1	19	Controller Port A1
Secondary	1B	1	20	Host 4, HBA 1 (HD1)

Table 3-4 Campus Configuration Connections (**continued**)

Site	Switch	Zone	Port	Connect To/From
Secondary	1B	1	21	Additional Host, HBA 1 (Optional)
Secondary	1B	1	22	Additional Host, HBA 1 (Optional)
Secondary	1B	1	23	Additional Host, HBA 1 (Optional)
Secondary	1B	2	24	Host 3, HBA 2 (HC2)
Secondary	1B	2	25	Controller Port B1
Secondary	1B	2	26	Host 4, HBA 2 (HD2)
Secondary	1B	2	27	Additional Host, HBA 2 (Optional)
Secondary	1B	2	28	Additional Host, HBA 2 (Optional)
Secondary	1B	2	29	Additional Host, HBA 2 (Optional)
Secondary	1B	3	30	Controller Port A2 (Dedicated RVM Port)
Secondary	1B	4	31	Controller Port B2 (Dedicated RVM Port)

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.
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 as presented in the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.

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).

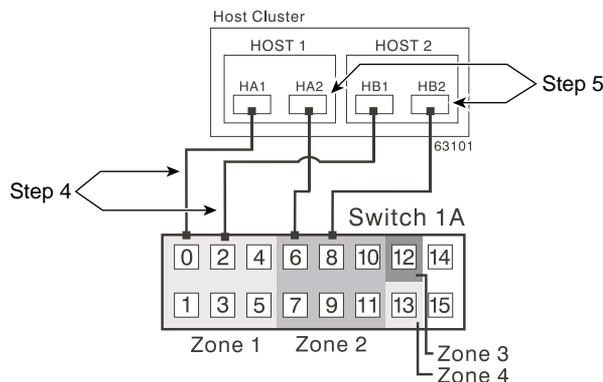


Figure 3-23 Host Bus Adapter Connections to Fibre Channel Switches

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 78.

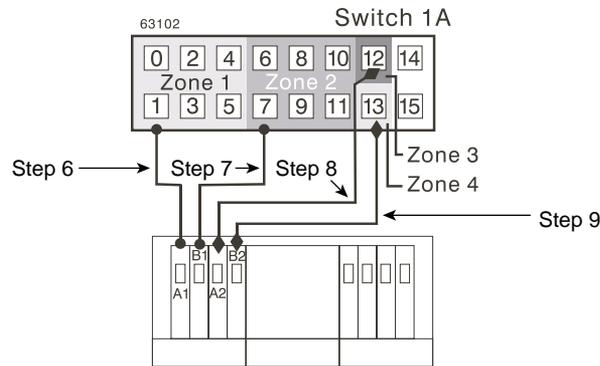


Figure 3-24 Storage Array Connections to Fibre Channel Switches

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).

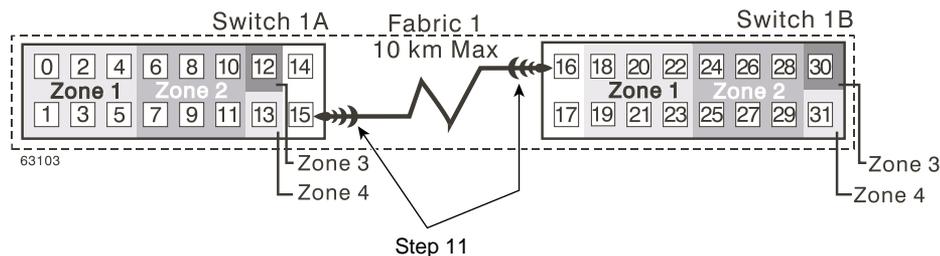


Figure 3-25 Connecting Switches over a Remote Network to Complete the Fabric Environment

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 121.

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.

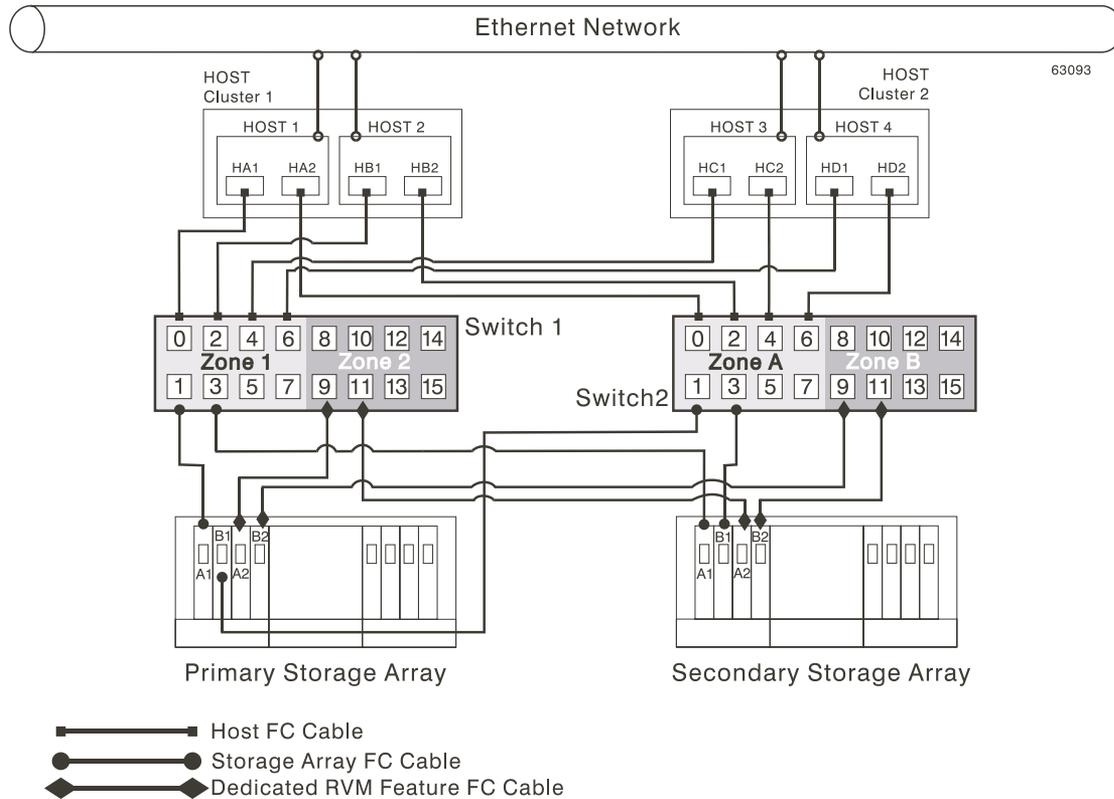


Figure 3-26 Intra-Site Configuration

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 95.

Cabling for the Intra-Site Configuration

Table 3-5 is provided for a convenient reference of port connections.

Table 3-5 Intra-Site Configuration Connections

Site	Switch	Zone	Port	Connect To/From
Primary	1	1	0	Host 1, HBA 1 (HA1)
Primary	1	1	1	Controller Port A1 (Primary Storage Array)
Primary	1	1	2	Host 2, HBA 1 (HB1)

Table 3-5 Intra-Site Configuration Connections **(continued)**

Site	Switch	Zone	Port	Connect To/From
Primary	1	1	3	Controller Port A1 (Secondary Storage Array)
Primary	1	1	4	Host 3, HBA 1 (HC1)
Primary	1	1	5	Additional Host, HBA 1 (Optional)
Primary	1	1	6	Host 4, HBA 1 (HD1)
Primary	1	1	7	Additional Host, HBA 1 (Optional)
Primary	1	2	8	Unused
Primary	1	2	9	Controller Port A2 (Primary Storage Array - Dedicated RVM Port)
Primary	1	2	10	Unused
Primary	1	2	11	Controller Port A2 (Secondary Storage Array - Dedicated RVM Port)
Primary	1	2	12	Unused
Primary	1	2	13	Unused
Primary	1	2	14	Unused
Primary	1	2	15	Unused
Secondary	2	A	0	Host 1, HBA 2 (HA2)
Secondary	2	A	1	Controller Port B1 (Primary Storage Array)
Secondary	2	A	2	Host 2, HBA 2 (HB2)
Secondary	2	A	3	Controller Port B1 (Secondary Storage Array)
Secondary	2	A	4	Host 3, HBA 2 (HC2)
Secondary	2	A	5	Additional Host, HBA 2 (Optional)
Secondary	2	A	6	Host 4, HBA 2 (HD2)
Secondary	2	A	7	Additional Host, HBA 2 (Optional)
Secondary	2	B	8	Unused
Secondary	2	B	9	Controller Port B2 (Primary Storage Array - Dedicated RVM Port)
Secondary	2	B	10	Unused

Table 3-5 Intra-Site Configuration Connections (**continued**)

Site	Switch	Zone	Port	Connect To/From
Secondary	2	B	11	Controller Port B2 (Secondary Storage Array - Dedicated RVM Port)
Secondary	2	B	12	Unused
Secondary	2	B	13	Unused
Secondary	2	B	14	Unused
Secondary	2	B	15	Unused

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.
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 as presented in the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.
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.

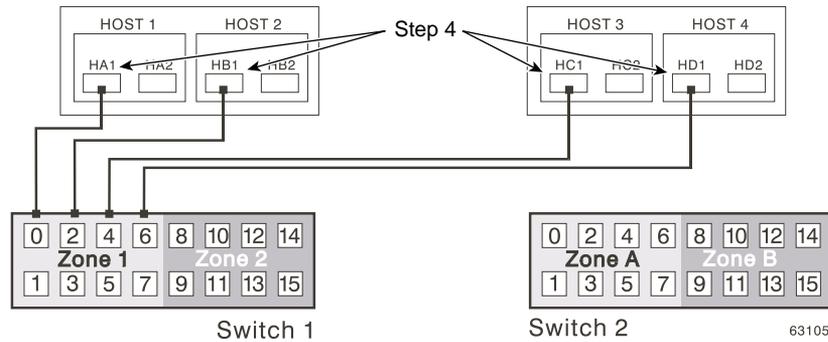


Figure 3-28 Primary Host Bus Adapter Connections to Fibre Channel Switches

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

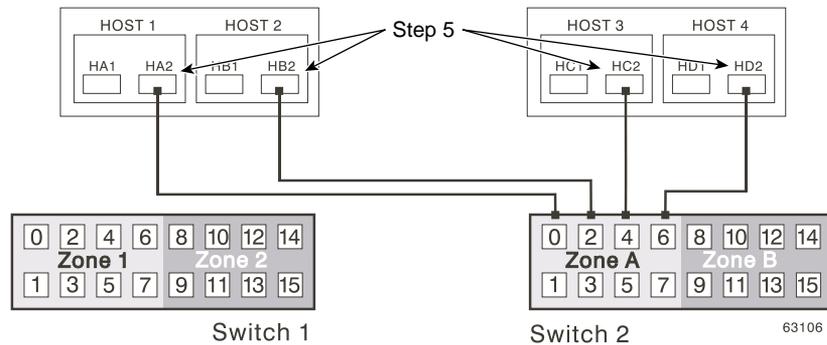


Figure 3-29 Secondary Host Bus Adapter Connections to Fibre Channel Switches

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 78.

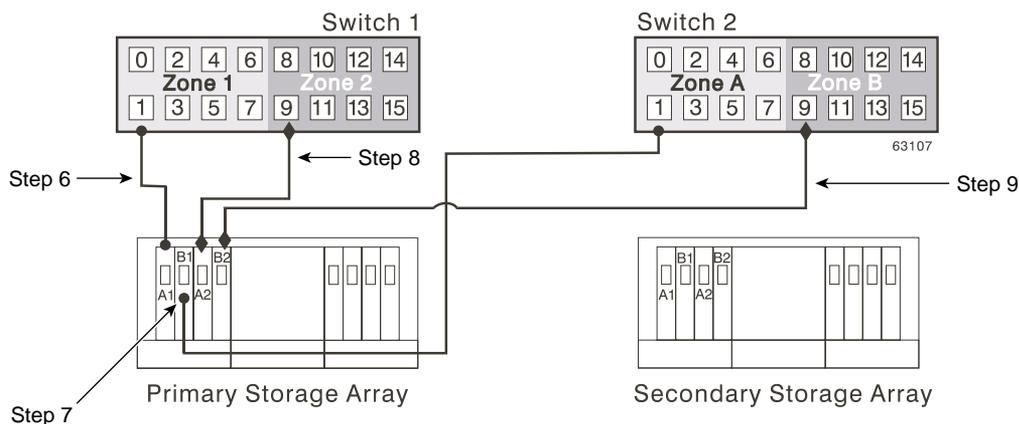


Figure 3-30 Primary Storage Array Connections to Fibre Channel Switches

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).

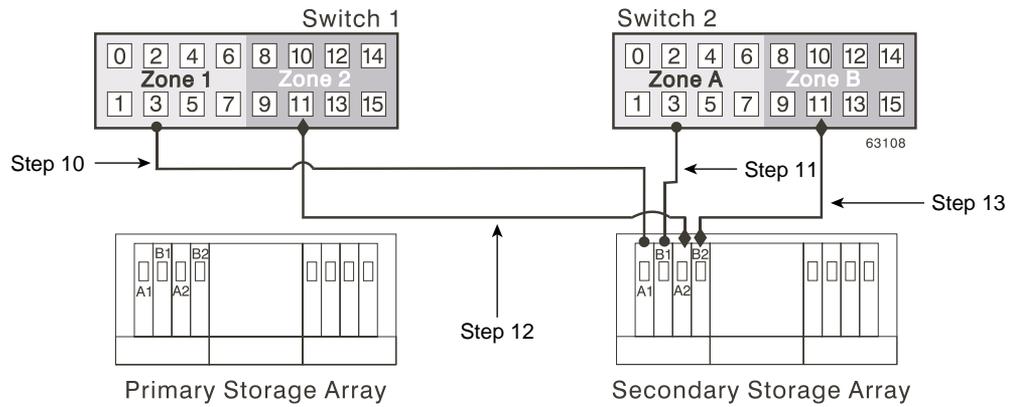


Figure 3-31 Secondary Storage Array Connections to Fibre Channel Switches

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 121.

Software Installation

The Remote Volume Mirroring premium feature is automatically included as part of the Total Performance Storage System Manager version 8.3x installation. However, the feature is not enabled or activated by default after installation.

Note: All storage arrays participating in mirror relationships must have firmware version 5.3x or later installed. For more information about upgrading firmware, see the *TP9500 Release Notes*.

For more information about installing Total Performance Storage System Manager version 8.3x and later, see the *SGI TP9400 and SGI TP9500 RAID Administration Guide*.

For more information about enabling or activating the Remote Volume Mirroring premium feature, see “Enabling and Activating the Remote Volume Mirroring Feature” on page 124.

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 TP9400 and SGI TP9500 RAID 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 145.

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 145.

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 8.3x or later. If this requirement is not met, see the *SGI TP9500 RAID Installation and Upgrade Guide* for storage management software upgrade procedures.
9. Select **OK**.
10. From the **Array Management** window, select **View > Storage Array Profile**.
11. Under the **All** tab, verify that the current firmware version is 5.3x or later. If this requirement is not met, see *TP9500 Release Notes* for firmware upgrade procedures.
12. Select **Close**.
13. 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 123.

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 124
- “Creating Mirror Relationships” on page 128
- “Maintaining Mirror Relationships” on page 131
- “Deactivating and Disabling the Remote Volume Mirroring Feature” on page 136

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, and 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.

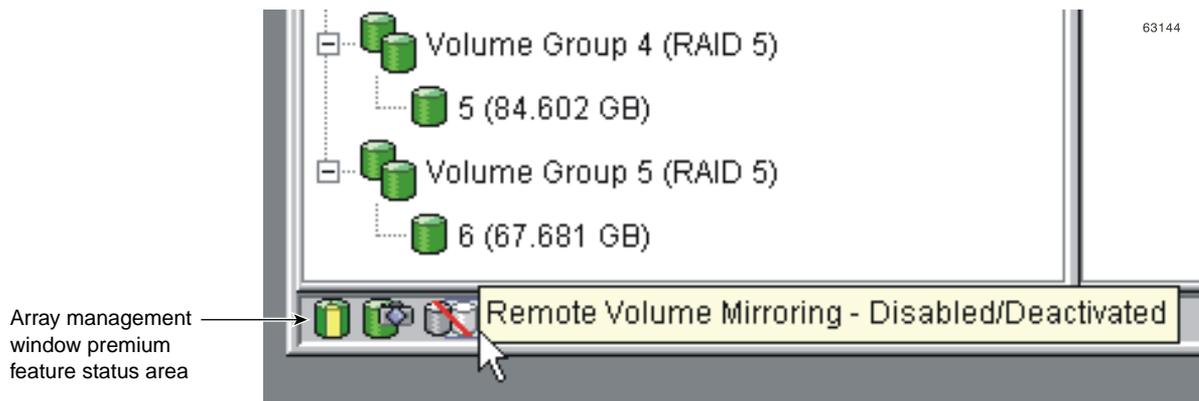


Figure 3-32 Remote Volume Mirroring Feature Status

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 126. To activate the feature, see “Activating the Remote Volume Mirroring Feature” on page 126.
- **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 126.
- **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 126.
- **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 128.

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 126.
 - Yes – Go to “Creating Mirror Relationships” on page 128.

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.

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.

- 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 126.
 - Yes – Go to “Creating Mirror Relationships” on page 128.

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.
- How to view mirror relationship configuration details.

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, or 5, and is of equal or greater size than 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 Priority Level” on page 133.

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.

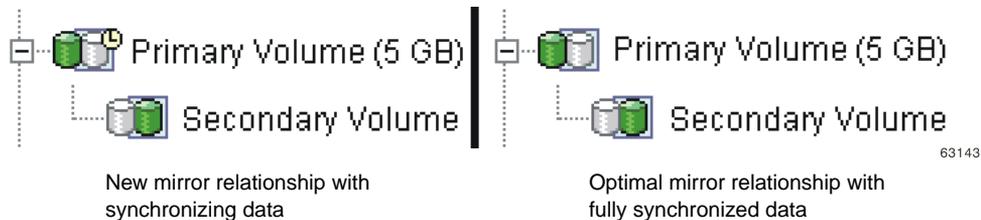


Figure 3-33 Synchronizing (Left) and Optimal (Right) Mirror Relationship

14. The next step is maintaining the mirror relationship. Go to “Maintaining Mirror Relationships” on page 131.

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 Mirror Relationship Details

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**.
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**.

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 Priority Level

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.

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**.

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 134.
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 137
- “Recovering from Switch Failures” on page 139
- “Recovering from Storage System Failures” on page 140
- “Reversing the Roles of Primary and Secondary Volumes” on page 143

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 Snapcopy Volume premium feature, which captures a point-in-time image of a volume. For more information on using snapcopy volumes, see “Snapcopy Volumes” on page 83 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 82 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 3-6 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 3-6 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.	__ 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.	__ Completed.
For more information on setting up Storage Partitioning, see the Array Management window online help.	
3. If possible, ensure that your host clustering software is configured to allow host failover between the primary storage array and secondary storage array participating in Remote Volume Mirroring.	__ 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. For information on how to directly connect hosts and storage arrays, see the *SGI TP9400 and SGI TP9500 RAID Owner's Guide*.
- 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 141.

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 142.

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 143.
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 82.
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 142.
 - No – Go to “Rebuilding a Damaged Site” on page 142.

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 94 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 82.
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 142.

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 131.

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 82.
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 143.

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 144.
- 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 144.

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 snapcopy volumes, demoting this volume to a secondary role will cause the associated snapcopy 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.

Troubleshooting topics include the following:

- “General Troubleshooting” on page 146
- “Inappropriate Fibre Channel Fabric Configurations” on page 153
- “Storage Array Component Failures” on page 154
- “General Switch Troubleshooting” on page 155

- “Remote Volume Mirroring Critical Events” on page 156
- “Controller Firmware Version 5.30 Upgrades” on page 158

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.

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 124 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 124 or the **Array Management** window online help.

- **When creating a Remote Volume Mirror, an error message listing the code `INVALID_MIRROR_CANDIDATE_VOL` 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 134 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 129 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 129.
- **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 129.
- **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 134 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 134 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 129 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 8 defined mirrors on an E4400 storage array and up to 32 defined mirrors on an E4600 or E5600 storage array.

For step-by step instructions, see “Removing Mirror Relationships” on page 134 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 134 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 129 or the **Array Management** window online help.

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 94.

- **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 94.

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 snapcopy volume created for a primary volume indicates that it has failed.**

Cause:

The base volume for the snapcopy 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 snapcopy volume on the secondary volume.

The base volume for a snapcopy volume cannot be a candidate for the secondary volume role in a Remote Volume Mirror, and the snapcopy 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 snapcopy volume will return to Optimal status.

For step-by-step instructions, see “Changing a Secondary Volume to a Primary Volume” on page 144 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 71.

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.

Controller Firmware Version 5.30 Upgrades

The Remote Volume Mirroring feature is designed so that it is not necessary for the primary storage array and the secondary storage to be running exactly the same controller firmware level. All inter-controller operations (including those needed for propagation of newly-written data from the primary volume to the secondary volume) have been implemented so as to support mixed firmware levels to the greatest extent possible. This allows firmware upgrades to be performed on both the primary and secondary storage arrays without any loss of data.

This subsection describes the procedure that should be followed for upgrading controller firmware (version 5.30 and later), to ensure that the Remote Volume Mirroring features works as intended. The preferred sequence for performing firmware upgrades on storage arrays participating in mirror relationships is detailed below.

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.

It should be noted that some configuration operations, such as establishment of new mirrors and modification of mirror parameters, may not be allowed between storage arrays that have different firmware levels, and will result in error messages being

received. However, the key data availability operations, including promotion from the secondary volume to the primary volume, will be allowed in such circumstances.

Index

A

access volume, 7, 8
Add Device option, 13
agent software
 see host-agent software
alert destinations, 46
alert notifications, 11
Array Management window
 described, 16
 failure notification in, 54
 functions, 16
 introduction to, 9
 monitoring Storage Arrays with, 46
Array Management window Help, xvii
Automatic Discovery option, 13
AVT (Auto-Volume Transfer) and RDAC, related to
 I/O, 30

B

battery age clock, 27
BOOTP or BOOTP-Compliant DHCP Server, 2

C

cache block size, 61
cache flush
 described, 28

 performance impacts, 28
 settings, 28
 start flush percentage, 61
 start percentage, 28
 stop flush percentage, 61
 stop percentage, 28
cache hit percentage
 defined, 61
 impact from cache flush settings, 28
cache read-ahead
 choosing a multiplier, 63
 defined, 63
cache settings, 66
Campus Configuration
 cabling, 110-114
 defined, 78
 redundancy characteristics, 90, 97
channel protection, 29
Components button, 18
configuration
 automatic, 66
 resetting a, 66
configuration file
 described, 11
 recovering a corrupted, 67
Contacting Device status, 49
controller, 36
 defined, 3
 listing components associated with, 68
 ownership, 66
 placing in active or passive mode, 68

- placing online or offline, 68
- viewing properties, 69
- controller host side I/O port
 - and read/write requests, 78
 - dedicated for mirroring, 78, 126, 128, 136
 - requirements, 78
 - zoning ports reserved for mirroring, 100
- controller module
 - see controller tray
- controller owner
 - and full synchronization, 72
 - and remote writes, 72
 - I/O path errors, 80
 - managing data replication, 75
 - of secondary volume, 72
 - prioritization of full synchronization, 81
 - read requests, 76
- controller resets, 73
- controller tray
 - batteries, 28
 - defined, 3
- copyback, 29
- Create Remote Volume Mirror Wizard, 129
- critical alerts, 14
- critical data, backing up, 137
- Critical Major Event Log (MEL), 156

- D**
- data protection strategies, 23
 - Hardware Redundancy, 27
- data replication
 - and link interruptions, 75
 - description of, 75
 - process illustration, 75
- Default Group, defined as a topological element, 36
- default LUN, 36
- default volume-to-LUN mapping
 - and volumes, 42
- Defined Mappings view, 36
- device
 - automatic discovery of, 64
 - correcting a partially managed, 67
 - manually adding, 64
 - removing, 64
- Device Table, 48
- Device Tree, 48
- directly managed storage array, 7
- Disabled/Activated status
 - changing the status, 136
 - defined, 125
- Disabled/Deactivated status
 - defined, 125
 - icon, 123
 - limitations of, 136
- disaster recovery
 - backing up critical data, 137
 - preparation steps, 137
- drive
 - defined, 3
 - failing, 67
 - initializing, 68
 - listing components associated with, 68
 - locating, 68
 - reconstructing, 68
 - reviving, 68
 - viewing properties, 69
- drive group
 - see volume group
- drive module
 - see drive tray
- drive tray
 - defined, 3
 - Hardware Installation and User Guides, xvii
- dual-primary condition, 144
- dual-secondary condition, 145

DVE (Dynamic Volume Expansion), also called
increase volume capacity, 33

E

email destinations, 67

emwdata.bin file, 11

Enabled/Activated status
changing the status, 136
defined, 125

Enabled/Deactivated status defined, 125

Enterprise Management window
figure, 12
functions, 10
introduced, 9
monitoring storage arrays with, 46
status icons displayed in, 48

Enterprise Management window Help, xvii

entire site failover
description of, 140

Event Log, 67

Event Monitor
setting up alert destinations, 51
synchronizing with Enterprise Management
window, 52

expected usage, of a volume, 32

F

failure notification
in the Array Management window, 54
types of, 54

Feature Key File, 126

Fibre Channel fabric configurations
defined, 78
distance limit, 80

Fibre Channel switches. See switches

firmware

downloading, 65

firmware version

for remote volume mirroring, 121
requirements, 87

firmware versions supported, 16

Fixing status, 48

free capacity, 4

adding to a volume group, 32, 65
consolidating on a volume group, 65
creating volumes from, 65
creating volumes on, 32
defined, 4
increasing storage array, 65

Free Capacity node, 4

full synchronization

after connectivity is restored, 77
after mirror relationship is created, 75
defined, 72
performance considerations, 81
time to complete, 81

H

hardware components, 1

Help

task reference, 64
using, xvii

heterogeneous hosts, overview and example, 43

Highest Availability Campus Configuration

cabling, 101-107
connection table, 101
defined, 78
redundancy characteristics, 88, 97

host

defined, 2

host bus adapter utility, 37

- host cluster software, 142
- host group, defined as a topological element, 36
- host port, defined as a topological element, 37
- host type, setting, 44
- host, defined as a topological element, 36
- host-agent managed Storage Array
 - described, 7
 - discovering newly attached, 64
- host-agent software
 - described, 8
 - where installed, 8
- hot spare drive
 - assigning, 65
 - configuring, 29
 - defined, 29
 - unassigning, 65

I

- increase volume capacity, also called Dynamic Volume Expansion (DVE), 33
- indicator lights, 68
- Installation and Support Guide, xvii
- Intra-Site Configuration
 - cabling, 121
 - defined, 79
 - redundancy characteristics, 91, 97, 114
- I/O access pattern, 62
- I/O data connection
 - see I/O data path
- I/O data path
 - described, 5
- I/O path error, 80
- I/O request rate
 - defined, 61
 - impact from cache flush settings, 28
 - modification priority effect on, 64

L

- link interruptions, 75
- logical unit number
 - see LUN
- Logical view, 48
 - status icon changes, 58, 59
- LUN
 - address space, 42
 - defined
 - using in a volume-to-LUN mapping, 42
 - see volume
- LUN (logical unit number), defined, 37

M

- Major Event Log (MEL), 156
- management domain
 - adding host-agent managed storage arrays to, 7
 - configuration file, 11
 - defined, 5
 - global e-mail alert settings, 67
 - populating, 13
 - use of Automatic Discovery, 13
- management station
 - see storage management station, 1
- Mappings View, overview, 34
- media scan
 - changing settings, 66
 - defined, 29
- medical imaging applications, 27
- MEL. See Major Event Log
- mirror relationships
 - maintaining, 131
 - maximum per storage array, 72
 - prerequisites for creating, 74
 - removing, 134, 150
 - viewing details, 132

mirror repository volume
 defined, 73
 deleting, 136
 information logged to, 75
 RAID level of, 74
 size of, 74
 missing volume, primary or secondary, 149
 modification operations, 63
 modification priority
 changing, 66
 described, 63
 effect on I/O request rate, 64
 setting, 64
 multimedia applications, 27
 multi-user environments, 27

N

Needs Attention status, 48
 network management connection
 combinations of, 7
 defined, 6
 host-agent, 8
 types of, 6
 Network Management Station
 defined, 2
 receipt of failure notifications, 55
 NVSRAM, downloading, 65

O

online Help
 see Help
 operating systems, list of certified, 88
 Optimal status, 48, 157
 orphan mirror, 150
 Overall Health Status, 48

P

password protection, 31, 68
 performance considerations, Remote Volume
 Mirrors, 81
 Performance Monitor, 61, 66
 Physical view
 changing tray order in, 65
 described, 18
 status icon changes, 58, 59
 power outage, 28
 preferred controller owner, 30
 primary volume
 and full synchronization, 72
 base volume for a snapcopy, 83
 controller ownership changes, 80
 creating, 129
 defined, 72
 ownership of, 79

R

RAID 0
 described, 26
 drive failure consequences, 26
 RAID 1
 described, 26
 drive failure consequences, 26
 performance, 62
 redundancy, 25
 RAID 3
 described, 27
 drive failure consequences, 27
 redundancy, 25
 RAID 5
 described, 27
 drive failure consequences, 27
 performance, 62

- redundancy, 25
- RAID level
 - and drive number constraints, 127
 - changing, 62, 66
 - described, 25
 - for secondary volume, 129, 130
 - of volume group for mirror repository volumes, 127
 - use on a volume group, 25
 - using an appropriate, 24
- RAID Module
 - see Storage Array
- read percentage, 62
- reconstruction, 29
- Recovery Guru
 - described, 56
 - identifying when to use, 67
 - Recovery Procedure, 57
 - Summary Area, 57
- redundancy
 - checking, 67
 - data, 24
 - of Fibre Channel Arbitrated Loops, 30
 - of Storage Array hardware, 23
- remaining capacity
 - see free capacity
- Remote Volume Mirror. See mirror relationships
- Remote Volume Mirroring feature
 - and other premium features, 82
 - firmware requirements, 87
 - hardware requirements, 87
 - introduction, 71
 - operating system requirements, 88
 - out of compliance, 151
 - software installation, 121
 - software requirements, 87
 - storage management software requirements, 87
 - troubleshooting, 145
 - verifying configuration before beginning, 122-123

- remote write, 75
- role reversal
 - with snapcopy volumes, 85
 - with Storage Partitioning, 83

S

- Script Editor, 11, 65
- secondary volume
 - and host reads and writes, 72
 - creating, 129
 - defined, 72
 - determining controller ownership, 79
 - failed, 77
 - minimum size, 129
 - offline, 77
 - Optimal status, 77
 - preparing for disaster recovery, 138
 - RAID level, 129, 130
 - restrictions with snapcopy premium feature, 84
- segment, 62
- segment size
 - changing, 62, 66
 - selecting an appropriate, 62
- site preparation, 94
- snapcopy repository volume, 83
- snapcopy volume
 - failure due to role reversal, 84-86
 - with Remote Volume Mirroring feature, 83
- SNMP destinations
 - specifying, 67
 - where stored, 11
- SNMP trap message
 - interpreting, 67
 - use for critical alerts, 2
- software components, 8
- specific volume-to-LUN mapping, 42
- start percentage, cache flush, 28

- status icon
 - changes during a recovery procedure, 59
 - changes during recovery procedure, 58
 - stop percentage, cache flush, 28
 - Storage Array
 - comment, 64
 - configuring, 23
 - connectivity, 5
 - defined, 2
 - directly managed, 7
 - host-agent managed, 7
 - indications of failure, 54
 - locating, 68
 - logical components, 4
 - management activities, 5
 - physical components, 3
 - renaming, 68
 - synchronizing clocks, 69
 - tuning options available, 61
 - viewing properties, 69
 - storage array component failures, 154
 - storage array health, summary of conditions, 46
 - Storage Array Management, 1
 - storage management software
 - where installed, 1
 - storage management software requirements, 87
 - storage management station
 - defined, 2
 - described, 1
 - Storage Partition
 - creating, 31, 65
 - defined, 4
 - major steps to creating, 40
 - Topology, 41
 - Storage Partitioning
 - overview and example, 38
 - topology, 36
 - Wizard, 38
 - storage partitioning
 - and secondary volumes, 83
 - switch zoning
 - and switch ports, 100, 109, 116
 - switches
 - cascading, 95
 - list of certified, 87
 - zoning, 95
 - Synchronization Priority Level
 - defined, 133
 - setting the, 81, 130
 - Synchronization-in-Progress status, 77, 78, 81
 - Synchronizing status, 157
 - synchronous mirroring, 76
- T**
- task reference, 64
 - terminology
 - functional equivalents, xv
 - text formatting conventions, xv
 - topological elements, defined, 36
 - Topology view, 36
 - transfer rate, 61
- U**
- unassigned capacity
 - see unconfigured capacity
 - unconfigured capacity
 - adding to a storage array, 65
 - creating volumes from, 65
 - creating volumes on, 32
 - defined, 4
 - Unconfigured Capacity Node, 4
 - Understanding guide
 - how to use, xv
 - organization, xv

prerequisite knowledge, xv

Unresponsive status

defined, 49

recovery, 67

Unsynchronized status, 157

after link interruption, 76

uplink port, 100, 109

UPS, 27

V

volume

and I/O accessibility, 72

candidates for mirroring, 74

creating, 31

defined, 4

initializing, 68

missing, 149

number of mirror relationships, 72

renaming, 68

sharing access to a, 42

viewing properties, 69

volume-to-LUN settings, 41

volume group

capacity of, 25

defined, 4

deleting, 65

initializing, 68

locating, 68

moving to another controller, 62

redistributing, 68

volume group data stripe, 62

volume modification priority

see modification priority

volume-to-LUN mapping

default, 42

defined, 38, 41

preference, 32

shared, 42

specific, 42

W

World Wide Name (WWN) updates, 151

write cache mirroring

described, 27, 28

how to enable, 27, 28

write caching

and data loss, 28

and performance, 28

described, 63

write caching without batteries, 27

WWN (World Wide Name), 37