

TPSSM7 RAID Software Concepts Guide for TP9400

007-4305-001

CONTRIBUTORS

Writing contributions by Olivier Clarinval

Edited by Connie Boltz

Production by Judi Holin

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

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

COPYRIGHT

© 2000 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 is a registered trademark and SGI and the SGI logo are trademarks of Silicon Graphics, Inc.

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 SYMplicity are trademarks of LSI Logic Corporation.

Record of Revision

Version	Description
001	August 2000 Initial Rev

Contents

Figures	vii
Tables	ix
About This Guide.	xi
Software and Hardware Terminology	xi
Additional Information	xii
Reader Comments	xiii
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	13
Enterprise Management Window Toolbar	14

- Array Management Window 14
 - Logical View and Physical View 15
 - Array Management Window Menu Options 17
 - Array Management Window Toolbar Options 18
- 2. Configuring Storage Arrays 19**
 - Strategies for Data Protection 19
 - Use an Appropriate RAID Level for Data Redundancy 19
 - Protect Data in the Controller Cache Memory 21
 - Configure Hot Spare Drives 22
 - Enable a Background Media Scan 23
 - Use Available Channel Protection When Possible 23
 - Configure Storage Array Password Protection. 24
 - Configure Alert Destinations for Critical Event Notification 24
 - Creating Volumes and Storage Partitions 24
 - Configuring Storage Array Capacity into Volumes 25
 - Introducing Storage Partitioning. 25
 - Example of Storage Partitions Configuration 30
 - Major Steps to Create Storage Partitions 32
 - Create Volumes on the Storage Array 32
 - Define Storage Partitions Topology. 33
 - Define Volume-to-LUN Mappings 34
 - Maintaining Storage Arrays in the Management Domain 37
 - Monitoring the Health of Storage Arrays 37
 - Customer Support Alert Notifications 42
 - Configuring the Customer Support Alert Notification 42
 - Recovering from Storage Array Problems 43
 - Tuning Storage Arrays 48
 - Task Reference 52
 - Index. 59**

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 and Array Management Windows	9
Figure 1-5	Enterprise Management Window	11
Figure 1-6	Device Tree with a Management Domain	12
Figure 1-7	Array Management Window	16
Figure 2-1	Storage Partitions Definitions Window	27
Figure 2-2	Storage Partitions Example	31
Figure 2-3	Defining Hosts and Host Groups in the Topology View	33
Figure 2-4	Defining a Volume-to-LUN Mapping	35
Figure 2-5	Host Group with Specific Mappings in the Topology View	36
Figure 2-6	Volume-to-LUN Mapping in the Mappings View	36
Figure 2-7	Monitoring the Health of Storage Arrays	38
Figure 2-8	Failure Notification in the Array Management Window	41
Figure 2-9	Recovery Guru Toolbar Button	43
Figure 2-10	Recovery Guru Window	45
Figure 2-11	Status Changes During Example Recovery Operation	46
Figure 2-12	Status Changes When an Example Recovery Operation is Completed	47

Tables

Table i	Terminology Change Reference	xiv
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	13
Table 1-6	Enterprise Management Window Toolbar Buttons.	14
Table 1-7	Array Management Window Menus.	17
Table 1-8	Array Management Window Toolbar Buttons	18
Table 2-1	RAID Level Configurations	20
Table 2-2	Creating Volumes from Unconfigured or Free Capacity	25
Table 2-3	Views in Storage Partitions Definitions Window	27
Table 2-4	Storage Partitions Terminology	28
Table 2-5	Storage Array Status Indications	39
Table 2-6	Tuning Options in Array Management Window	49
Table 2-7	Task Reference: Populating a Management Domain	52
Table 2-8	Task Reference: Storage Array Configuration	53
Table 2-9	Task Reference: Performance and Tuning	54
Table 2-10	Task Reference: Event Notification	55
Table 2-11	Task Reference: Recovering From Problems	55
Table 2-12	Task Reference: Miscellaneous System Administration	56

About This Guide

This guide explains the terminology and features of the TPSSM7 storage management software for the TP9400. 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 TPSSM7 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.

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.

Additional Information

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

- *SGI TP9400 RAID Owner's Guide* (007-4304-001) 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 RAID IRIX Administration Guide* (007-4306-001) 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 *SGI TP9400 RAID IRIX 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 RAID IRIX Administration Guide* for information specific to the IRIX operating environment.

To obtain SGI documentation, see the SGI Technical Publications Library at <http://techpubs.sgi.com>.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number can be found on the back cover.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
`techpubs@sgi.com`

- Use the Feedback option on the Web page of the Technical Publications Library at:
<http://techpubs.sgi.com>
- 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.

SGI values your comments and will respond to them promptly.

Storage Array Management Architecture and Software Components

This chapter describes the major hardware components in the TPSSM7 storage management architecture and the physical and logical components of a storage array. This section 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 the host, rather than on a storage management station. From the host, storage management commands are sent over the I/O data connection, not over the network. However, the TCP/IP networking protocol is still used for the communication between the client software and the controllers.

Major Hardware Components

Table 1-1 lists the major hardware components that are involved in the TPSSM7 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	<p>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).</p> <p>See “Physical Components of a Storage Array” on page 3 and “Logical Components of a Storage Array” on page 4 for more information.</p>
Storage management station	Computer you use to manage storage arrays on the network.
Network management station (optional)	<p>Console with installed SNMP-compliant network management software that receives and processes information about managed network devices using SNMP (Simple Network Messaging Protocol).</p> <p>Important: The storage management software sends critical alerts (using SNMP trap messages) to an NMS destination that you specify using the Edit >> Alert>> Destinations option in the Enterprise Management window. For the critical alerts to be sent, the Enterprise Management software must be running (the window can be minimized).</p>
BOOTP or BOOTP-compliant DHCP server	Used to assign the network-specific information such as 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.

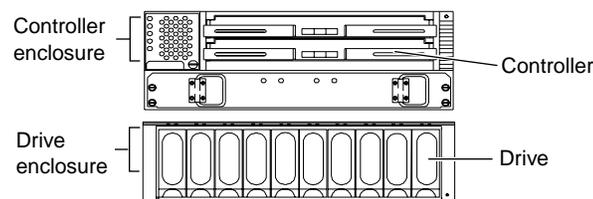


Figure 1-1 Physical Components of a Storage Array

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 1-7 on page 16). 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 “Introducing Storage Partitioning” on page 25.

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 “Introducing Storage Partitioning” on page 25).
- 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.

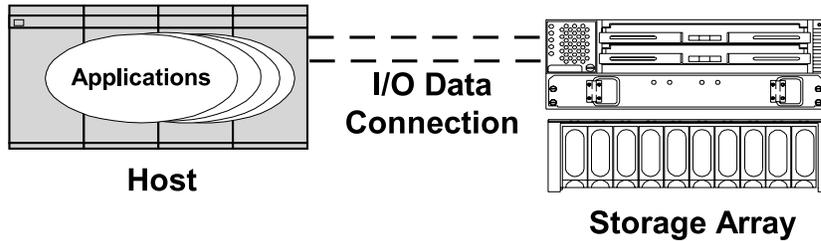


Figure 1-2 I/O Data Connection

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.

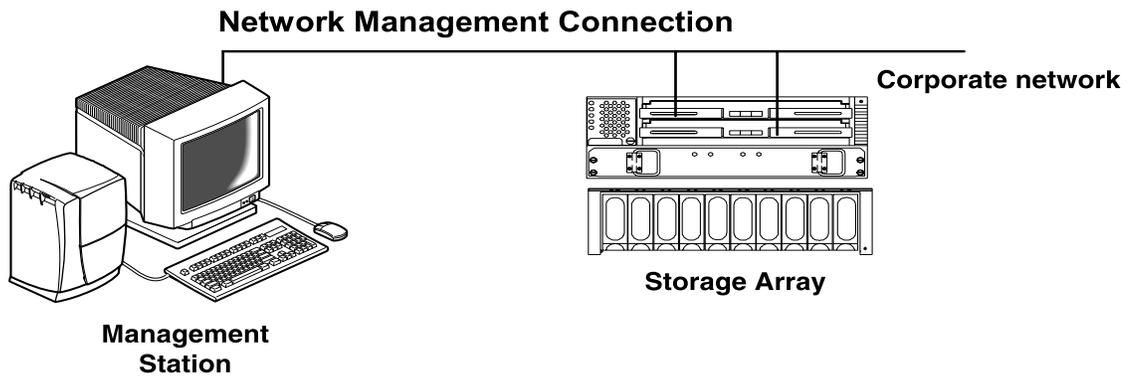


Figure 1-3 Network Management Connections

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 as "out-of-band" management.

For more information about directly managed storage arrays, see the *SGI TP9400 RAID IRIX 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 TPSSM7 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 #31. See Table 1-4 on page 8 for a description of the host-agent software. This form of RAID management is also referred 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 "Additional Information" on page xii.

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

Major Software Components

The TPSSM7 storage management software consists of three major components, described in Table 1-4.

Table 1-4 Major Software Components

Component	Description
TPSSM7 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 RAID IRIX administration Guide</i> for details.</p> <p>For more information, see “Client Software Windows” on page 9.</p>
TPSSM7 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 access volume. This volume uses one of the logical unit numbers (LUNs) available on the host.</p> <p>In addition, the host-agent package contains the SM7devices utility that correlates the volumes you create using the storage management software with their operating system device names. 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 RAID IRIX Administration Guide</i> for details.</p>

Client Software Windows

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

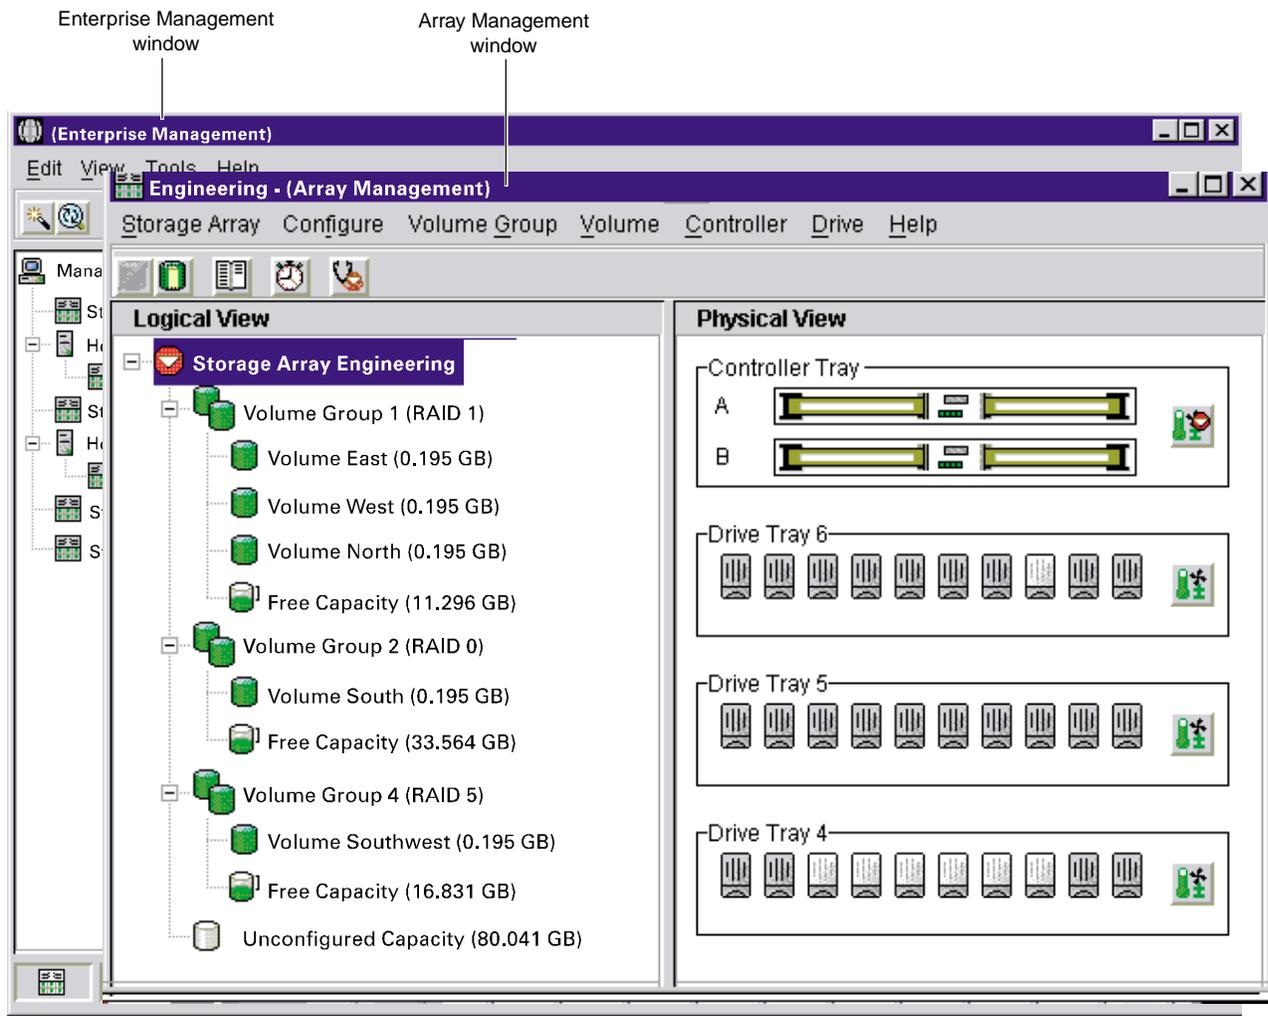


Figure 1-4 Enterprise Management and Array Management Windows

Enterprise Management Window

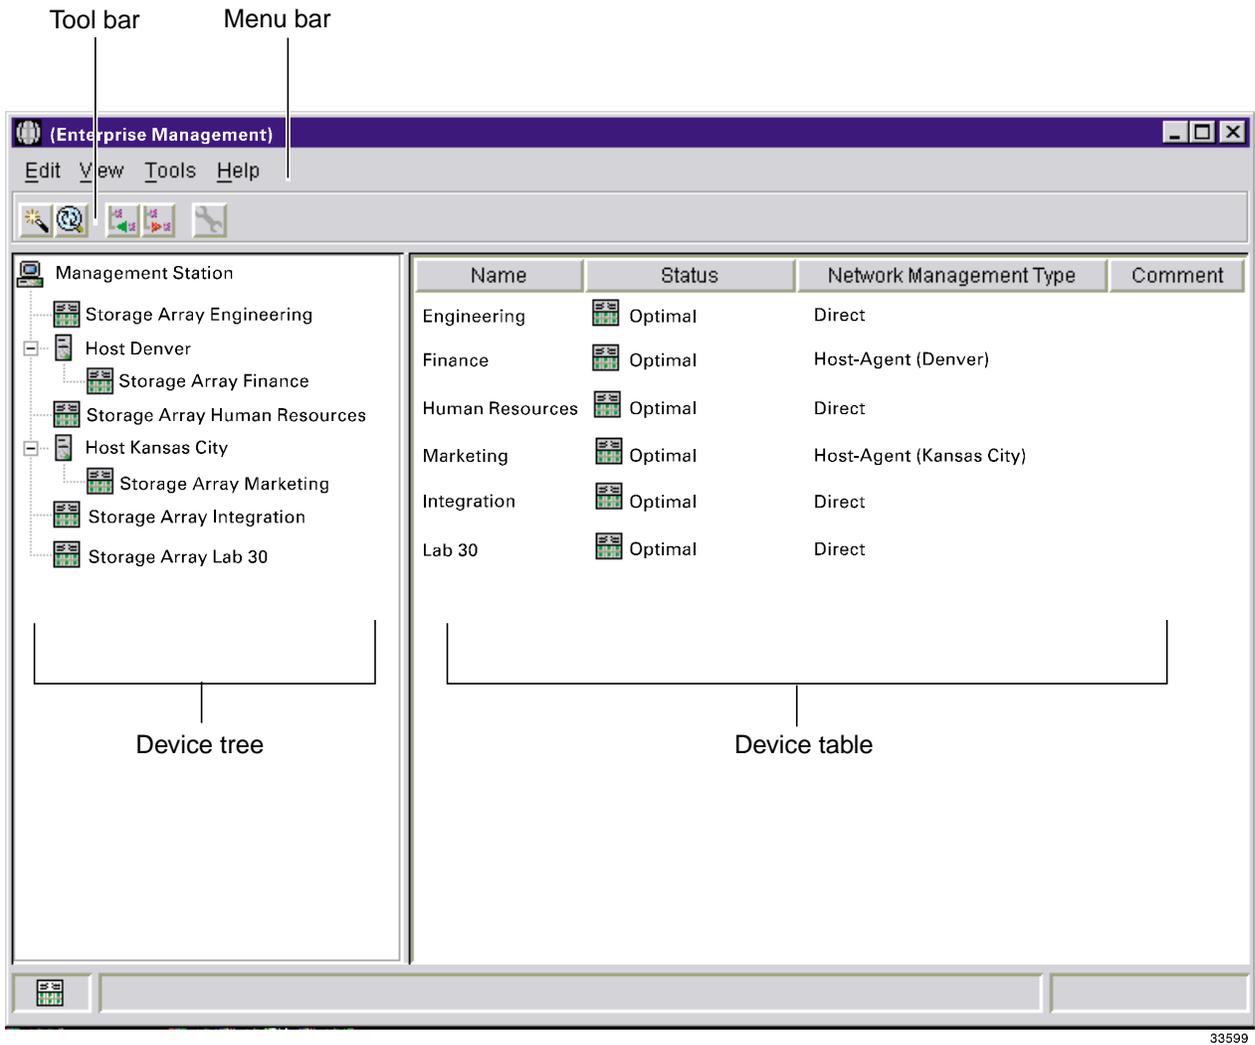
The Enterprise Management window, shown in Figure 1-5, is the first window to appear when you start the software. You use the Enterprise Management window to:

- Add and discover the storage arrays you want to manage (see “Additional Information” on page xii).
- Provide a comprehensive view of all storage arrays in the management domain.
- Perform batch storage array management tasks using the Script Editor.
- Configure destinations (email 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 email/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 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 37.



33599

Figure 1-5 Enterprise Management Window

Device Tree

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

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

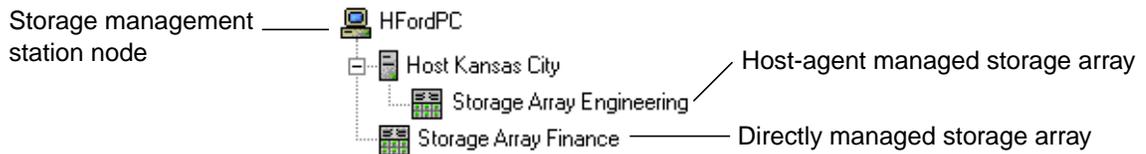


Figure 1-6 Device Tree with a Management Domain

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

1. Use the **Automatic Discovery** option to detect directly managed and host-agent managed storage arrays on the local sub-network 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. 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 1-5.)

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.
View	Allows you to sort the entries in the Device Table by Name, Status, Management Type, or Comment. This menu also shows Partially Managed Devices , 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 sub-network or rescan to find storage arrays newly attached to a host. You can also open an Array Management window to manage a selected storage array and open the Script Editor to perform batch management tasks.
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 discover new devices 	Activates the Automatic Discovery option that discovers hosts and storage arrays on the local sub-network 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 only available when you select a host in the Device Tree.
Add host/device 	Opens the Add Device dialog for manually adding hosts or controllers that comprise a storage array 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.
Launch an Array Management window 	Opens an Array Management window for the selected storage array. Note: This option is only available when you select a storage array in the Device Tree or Device Table.

See Figure 1-5 on page 11 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

You use the Array Management window to configure and maintain the logical and physical components comprising a storage array. See Figure 1-7 on page 16.

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 other storage arrays simultaneously.

Logical View and Physical View

The **Logical View** in the left pane shows the organization of storage array capacity into volume groups and volumes.

The **Physical View** in the right pane shows the physical devices in the storage array, such as controllers, drives, and other components. Selection of a volume or other entity in the Logical View shows you the associated physical components in the Physical View.

A Components button in every enclosure presents the status of all enclosure components and shows the temperature status of the enclosure.

These elements are labeled in Figure 1-7 on page 16. For more specific information about the Logical and Physical Views and the Components button, see the Array Management window Help.

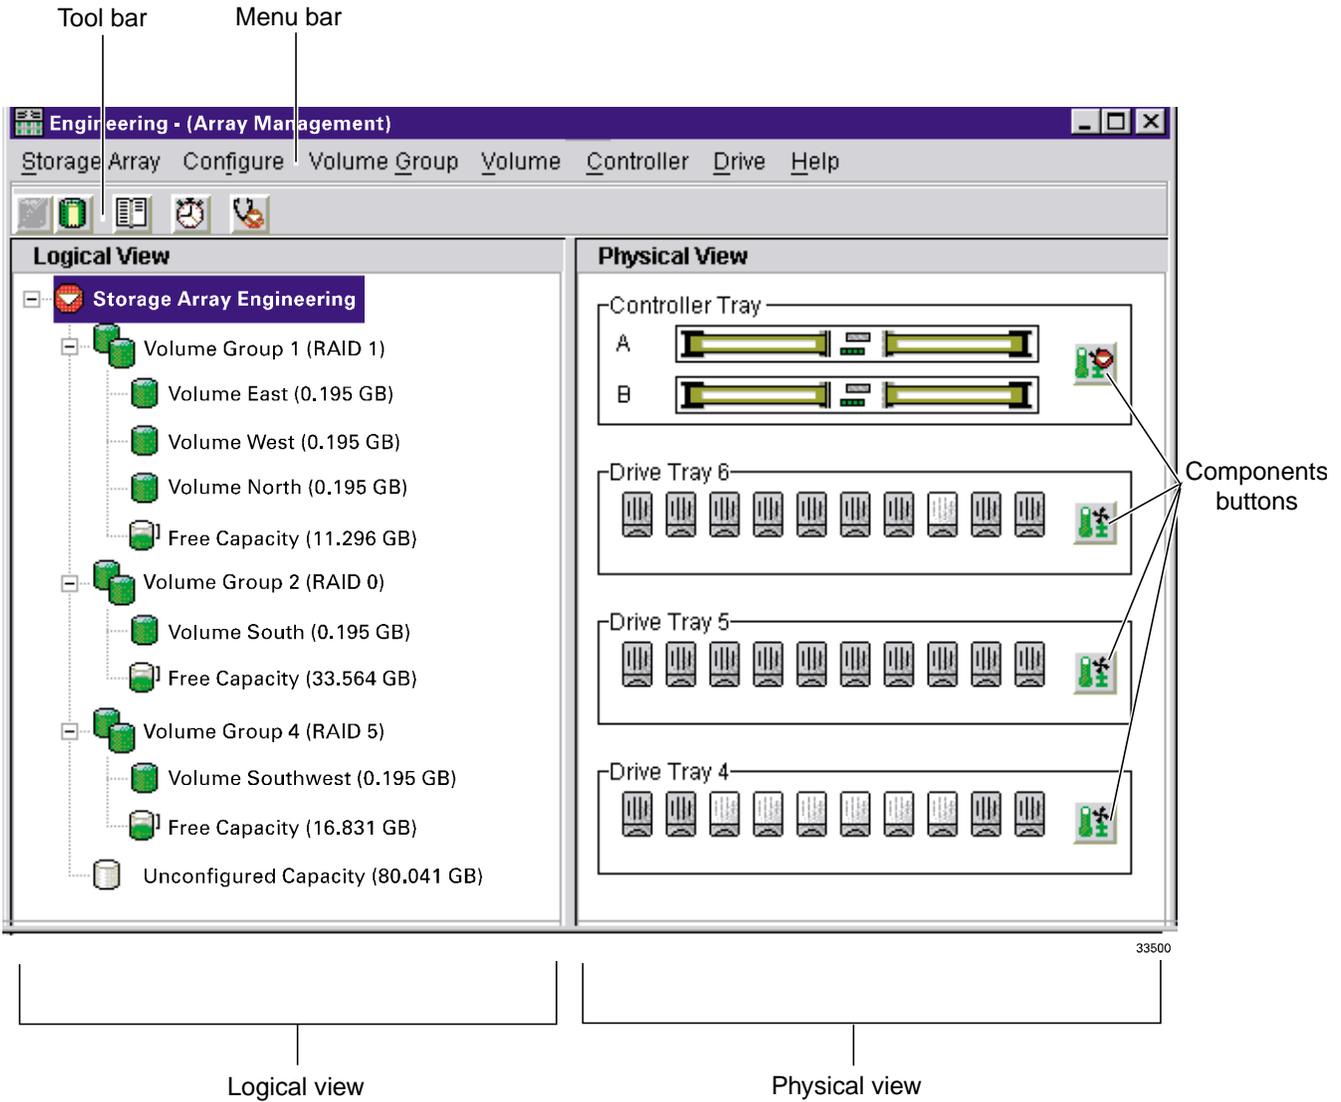


Figure 1-7 Array Management Window

Array Management Window Menu Options

The Array Management window menus on the menu bar are described in Table 1-7.

Table 1-7 Array Management Window Menus

Menu	Description
Storage Array	Allows you to perform storage array management operations such as monitoring performance, changing cache settings, starting Recovery Guru, or viewing detailed profile information about the storage array.
Configure	Allows you to configure the logical components of a storage array (volumes, volume groups, and storage partitions) from the physical capacity. You can also change the order of the enclosures in the Physical View to match the actual layout of the physical components.
Volume Group	Allows you to perform volume group storage management operations such as changing the RAID level, changing controller ownership, or capacity expansion.
Volume	Allows you to perform volume storage management operations such as changing segment size, initializing a volume, or viewing and setting volume properties.
Controller	Allows you to perform controller storage management operations such as placing a controller online or offline, or viewing or setting controller properties.
Drive	Allows you to perform drive storage management operations such as assigning or unassigning a hot spare, or viewing drive properties.
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 described in Table 1-8.

Table 1-8 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.
Configure storage partitioning 	Allows you to define the topological elements attached to your storage array and define volume-to-LUN mappings to create storage partitions. For more on storage partitions, see “Introducing Storage Partitioning” on page 25.
View diagnostic event log 	Opens the Event Log Viewer. Note: This log is for customer support use.
Monitor performance 	Opens the Performance Monitor.
Recover from failures 	Opens the Recovery Guru. Note: If the storage array is in the Needs Attention state,  the Recovery Guru toolbar button flashes.

See Figure 1-7 on page 16 for an example of the Array Management window toolbar. For more information on using the toolbar options, see the Array Management window Help.

Configuring Storage Arrays

This section 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.

Strategies for Data Protection

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. You can also use the software to implement the following data protection strategies:

- “Use an Appropriate RAID Level for Data Redundancy” on page 19
- “Protect Data in the Controller Cache Memory” on page 21
- “Configure Hot Spare Drives” on page 22
- “Enable a Background Media Scan” on page 23
- “Use Available Channel Protection When Possible” on page 23
- “Configure Storage Array Password Protection” on page 24
- “Configure Alert Destinations for Critical Event Notification” on page 24

Use an Appropriate RAID Level for Data Redundancy

RAID is an acronym for Redundant Array of Independent Disks. It describes a storage solution in which part of the storage capacity is used to store redundant information about user data stored on the remainders of the storage capacity.

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 level 0, 1, 3, and 5. Table 2-1

on page 20 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. You can use the redundancy data can be used to quickly reconstruct information on a replacement drive if a drive fails.

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 (continued) RAID Level Configurations

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>

Protect Data in the Controller Cache Memory

Write caching can increase 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**, 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. 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 49), and therefore the I/O request rate (also described in Table 2-6 on page 49). 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 enclosure protect the cache against power outages. Change the controller enclosure batteries at the recommended time intervals. The storage management software features a battery age clock you can set when you replace a battery. This clock will keep track of the age (in days) of the battery so you know when it is time to replace the battery. You will receive critical alert notification when the battery is nearing expiration and when it has reached expiration.

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 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 does not have working batteries.

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

Configure Hot Spare Drives

You can 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**.

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**.

Enable a Background Media Scan

A **media scan** is a background process that runs on all volumes in the storage array for which it has been enabled and provides error detection on the drive media. 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

When you create volumes from unconfigured capacity, volume group candidates are shown on the Create Volume - Select Volume Group dialog with information about whether or not the volume group candidate has channel protection. If the interface between the controller and the drives is SCSI, 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 were present on a single drive channel. If the interface between the controller and drives is Fibre Channel, 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.

Configure Storage Array Password Protection

You can configure a password for each storage array you manage for added security. The password is stored on the storage array. Therefore, if desired, you need to provide a password for each storage array.

When you start an Array Management window for a particular storage array, you are prompted for a password the first time you attempt a destructive operation. You are asked for the password only once during a single management session.

Configure Alert Destinations for Critical Event Notification

There are flexible options available for configuring alert notification destinations (email and/or SNMP). You can set up alert notification destination addresses for notifications about the following:

- Every storage array in the management domain
- Every storage array attached and managed through a particular host
- An individual storage array

This software also enables you to validate potential destination addresses and, for a management domain, specify global email alert settings for mail server and sender email address.

Note: The Enterprise Management window must be open to receive alert notifications about storage arrays in the management domain.

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 “Additional Information” on page xii.</p>
Free Capacity	<p>Because free capacity only resides 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>

Introducing Storage Partitioning

A storage partition is a logical entity consisting of one or more storage array volumes that can be accessed by a single host or shared among hosts that are part of a host group. You create a storage partition by defining a collection of hosts (a host group) or a single host,

and then defining a volume-to-logical unit number (LUN) mapping. This mapping allows you to define which host group or host will have access to a particular volume in the storage array.

Storage partitioning is a premium feature of the software that must be configured by either you or your storage vendor supplier before it can be used. There is a maximum number of storage partitions that can be created on the storage array depending on the features that have been enabled. The benefits of storage partitioning include ease of management, amortization of costs, scalability, and flexibility.

Use the Configure >> Storage Partitioning option in the Array Management window to display the Storage Partitioning Definitions window. This window provides all of the options necessary to create storage partitions.

The Storage Partitioning Definitions window has two views, Topology and Mappings, shown in Figure 2-1 and described in Table 2-3. The other terms in Figure 2-1 are defined in Table 2-4 on page 28.

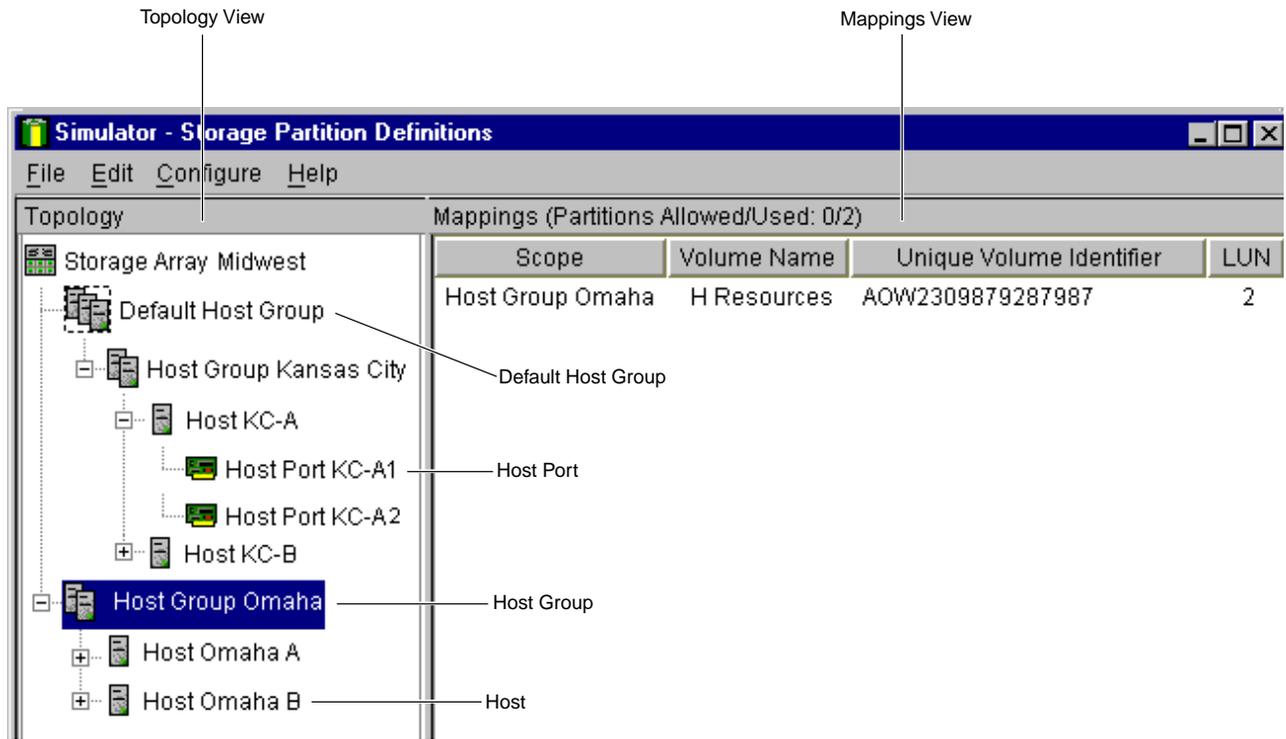


Figure 2-1 Storage Partitions Definitions Window

Table 2-3 Views in Storage Partitions Definitions Window

View	Description
Topology	Shows defined topological elements (host groups, hosts, host ports) and the default host group.
Mappings	Displays the default and/or specific volume-to-LUN mappings in the storage array (in table form). Volumes information includes: topological entity that can access the volume, volume name, unique volume identifier, and LUN number used to access the volume.

The terminology listed in Table 2-4 is used in the Storage Partitions Definitions window. (See Figure 2-1 on page 27 as you are reading these terms.)

Table 2-4 Storage Partitions Terminology

Term	Description
Storage partition topology	<p>Collection of topological elements (default host group, host group, host, and host port) shown as nodes in the Topology View of the Storage Partitions window.</p> <p>Note: You must define the various topological elements if you want to define specific volume-to-LUN mappings for host groups and /or hosts. Even if you are not planning to use storage partitions, you can still define hosts and host ports that are connected to the particular storage array and view the default volume-to-LUN mappings to determine which logical unit numbers are being used to access the volumes in the storage array.</p>
Default host group	<p>Standard node in the Topology View that designates all host groups, hosts, and host ports as follows: (1) they do not have any specific volume-to-LUN mappings, and (2) they share access to any volumes that were automatically assigned default LUNs by the controller firmware during volume creation.</p>
Host group	<p>Optional topological element that you define if you want to designate a collection of hosts that will share access to the same volumes.</p>
Host	<p>Computer that is attached to the storage array and accesses various volumes on the storage array through its host ports. You can define a host as a separate entity or as part of a host group. You can define specific volume-to-LUN mappings to an individual host as well as designate the host to be part of a host group that shares access to one or more volumes. Hosts can only be within one parent level host group at any time.</p>

Table 2-4 (continued) Storage Partitions Terminology

Term	Description
Host Port	<p>Physical connection on the host adapter that resides within a host. When the host adapter has only one physical connection (host port), the terms host port and host adapter are synonymous.</p> <p>The host ports are automatically discovered by the software. A host port is the actual physical connection that allows a host to gain access to the volumes in the storage array. Therefore, if you want to define specific volume-to-LUN mappings for a particular host and create storage partitions, you must define its associated host ports.</p> <p>Initially, all discovered host ports belong to the default host group and have access to any volumes that were automatically assigned default LUNs by the controller firmware during volume creation.</p> <p>Important: Use the host adapter boot-up utility to find the worldwide name of the host port. This is the host port identifier shown in the Define New Host Port dialog in the Storage Partitions Definitions window. Create a list of the worldwide names for host ports on a particular host and their defined host port names in the Topology View.</p>
Logical unit number (LUN)	<p>The number a host uses to access a volume. Each host has its own LUN address space. Therefore, different hosts may use the same LUN to access different volumes. However, a volume can only be mapped to a single LUN and a single host group or host.</p> <p>For example, Host 1 may access volume Engineering by using LUN 5 and Host 2 may also access volume Financial by using LUN 5 (see Figure 2-2 on page 31).</p>

Table 2-4 (continued) Storage Partitions Terminology

Term	Description
Default volume-to-LUN mapping	<p>During volume creation, you can specify whether you want to have a default LUN mapping automatically assigned to the volume. When a host group or host is initially created, it is placed within the default host group, and can access all of the default volume-to-LUN mappings. A host group or host can access either volumes with default mappings or specific mappings, but not both.</p> <p>Note: Any volumes that were created using version 6.x of this software are automatically given default volume-to-LUN mappings. If desired, you can change these mappings using the Change Mapping option.</p>
Specific volume-to-LUN mapping	<p>Occurs when you select a defined host group or host in the Topology View and select the Define New Mapping option to assign a specific volume and logical unit number (volume-to-LUN mapping). This designates that only the selected host group or host can access the particular volume through the assigned LUN. You can define one or more specific volume-to-LUN mappings for a host group or host.</p> <p>Note: You must enable the Storage Partitioning feature in order to create specific mappings.</p>

Example of Storage Partitions Configuration

In the example shown in Figure 2-2 on page 31, four hosts (Omaha A, Omaha B, KC-A, and KC-B) are connected to storage array Midwest. In our example, you want Omaha A and Omaha B to have access to volumes Marketing and HResources, and you want KC-A and KC-B to have access to volumes Legal and Engineering. You also want host KC-B to have exclusive access to volume Financial. Therefore, you must configure three storage partitions.

The first partition consists 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 with the host group Kansas City. These volumes are accessed by hosts KC-A and KC-B in host group Kansas City by using LUNs 2 and 4, respectively.

The third partition consists of volumes Marketing and HResources. 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 by using LUNs 7 and 2, respectively.

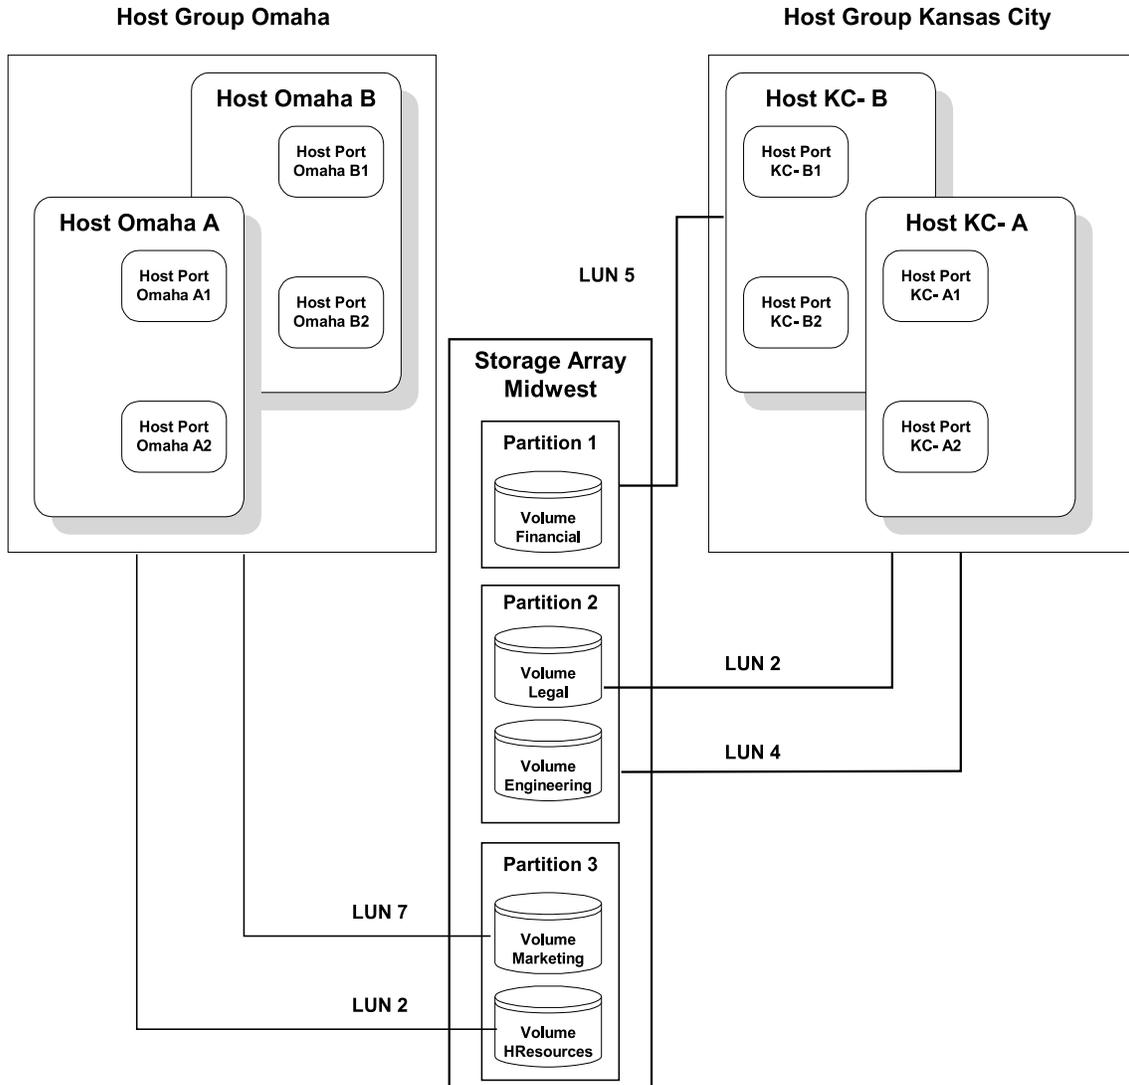


Figure 2-2 Storage Partitions Example

The next section explains how this example was configured.

Major Steps to Create Storage Partitions

Creating storage partitions has the following three major steps, which are detailed in the sections that follow.

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

During volume creation, select either of the following two settings for volume-to-LUN mapping:

- **Default.** This setting specifies that a LUN should be automatically assigned to the volume. It also allows host groups or hosts that do not have specific volume-to-LUN mappings (designated by the default host group node in the Topology View) to have access to the volume. Specify this setting if you will not use storage partitions.
- OR-
- **None (map later).** This setting specifies that a LUN will not be assigned to the volume. It also allows you to define a specific volume-to-LUN mapping and to create storage partitions using the Configure >> Storage Partitioning option. Specify this setting if you will use storage partitions. In Figure 2-2 on page 31, the volumes in storage array Midwest were given this setting. Then after the storage partitions topology was defined (step 2), specific volume-to-LUN mappings were assigned (step 3).

Define Storage Partitions Topology

Use the Storage Partitioning Definitions window (see Figure 2-1 on page 27) to define the host groups, hosts, and associated host ports that will have specific volume-to-LUN mappings.

The topological elements that you define (host groups, hosts, and host ports) and the default host group are displayed in the Topology View, as shown in the following Figure 2-3.

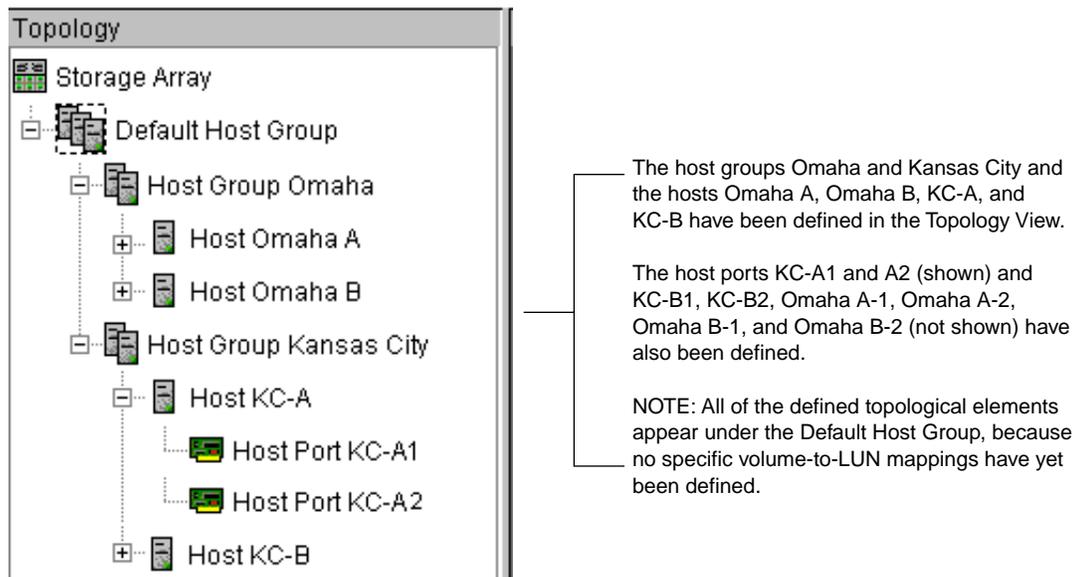


Figure 2-3 Defining Hosts and Host Groups in the Topology View

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

- Volumes with default volume-to-LUN mappings. In this case, the host group or host is part of the default host group.
- Volumes for which it has been granted access through a specific volume-to-LUN mapping.

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).

If you define specific volume-to-LUN mappings (as explained in the following section) 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.

In Figure 2-2 on page 31, all host groups and hosts have specific mappings and can access only the volumes to which they have been granted access through a volume-to-LUN mapping.

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 according to the following general steps:

1. Select a defined host group or host in the Topology View.
2. Select the Define New Mappings option.
3. Select a volume that you want the host group or host to be able to access.
4. Select the LUN that you want the host group or host to use to access the volume.

In Figure 2-4 on page 35, a new volume-to-LUN mapping is defined for volume HResources, which allows host group Omaha to access this volume through LUN 2, as shown in Figure 2-2 on page 31.

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.

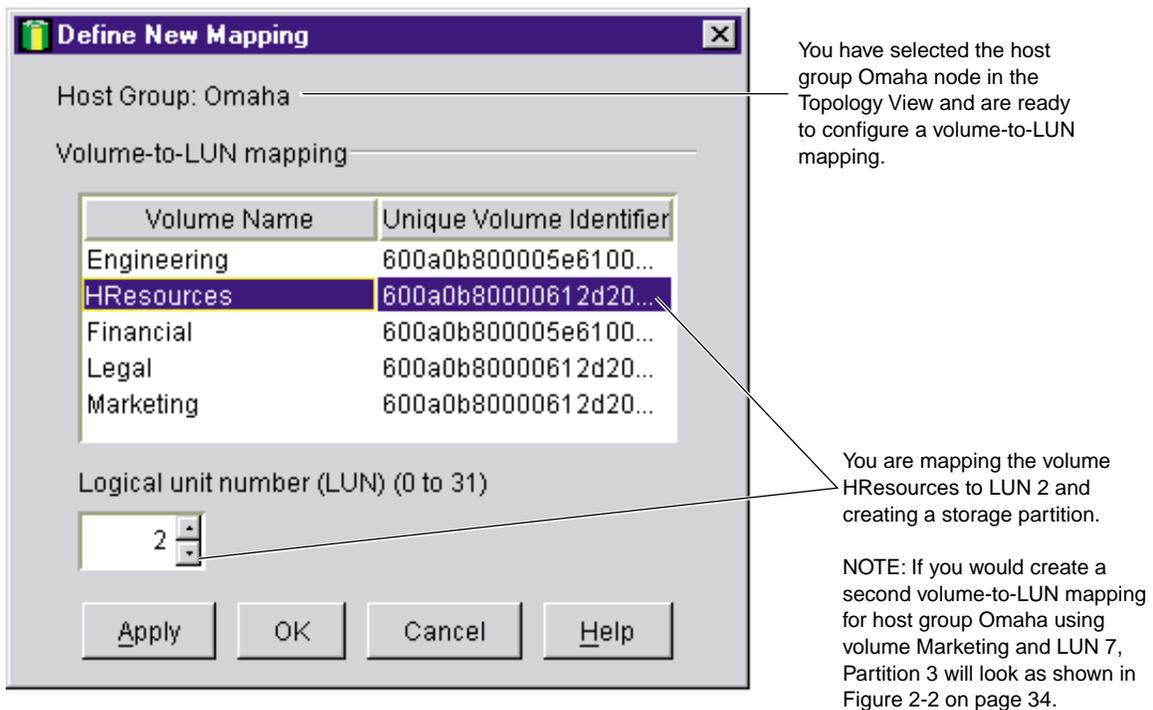
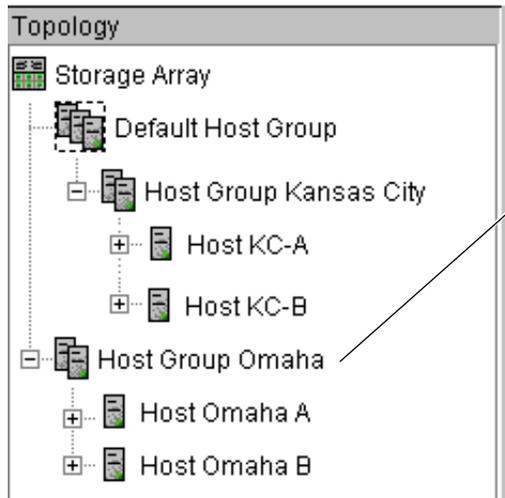


Figure 2-4 Defining a Volume-to-LUN Mapping

By defining the specific volume-to-LUN mapping for volume HResources using LUN 2, the following results occur in the Topology View (shown in Figure 2-5) and in the Mappings View (shown in Figure 2-6 on page 36).



After you have defined the first specific volume-to-LUN mapping (volume HResources using LUN 2) for host group Omaha, the host group and its associated hosts and host ports move out from under the Default Host Group and a storage partition is created.

The associated hosts (Omaha A and Omaha B) can now access volume HResources. However, since they are no longer part of the Default Host Group they cannot access any volumes with default volume-to-LUN mappings.

Figure 2-5 Host Group with Specific Mappings in the Topology View

Figure 2-6 shows the result in the Mappings View when you select the host group Omaha in the Topology View.

Mappings			
Scope	Volume Name	Unique Volume Identifier	LUN
Host Group Omaha	H Resources	A0W2309879287987	2

Figure 2-6 Volume-to-LUN Mapping in the Mappings View

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 destinations for critical event notification (see “Configuring the Customer Support Alert Notification” on page 42 or the Enterprise Management window Help).

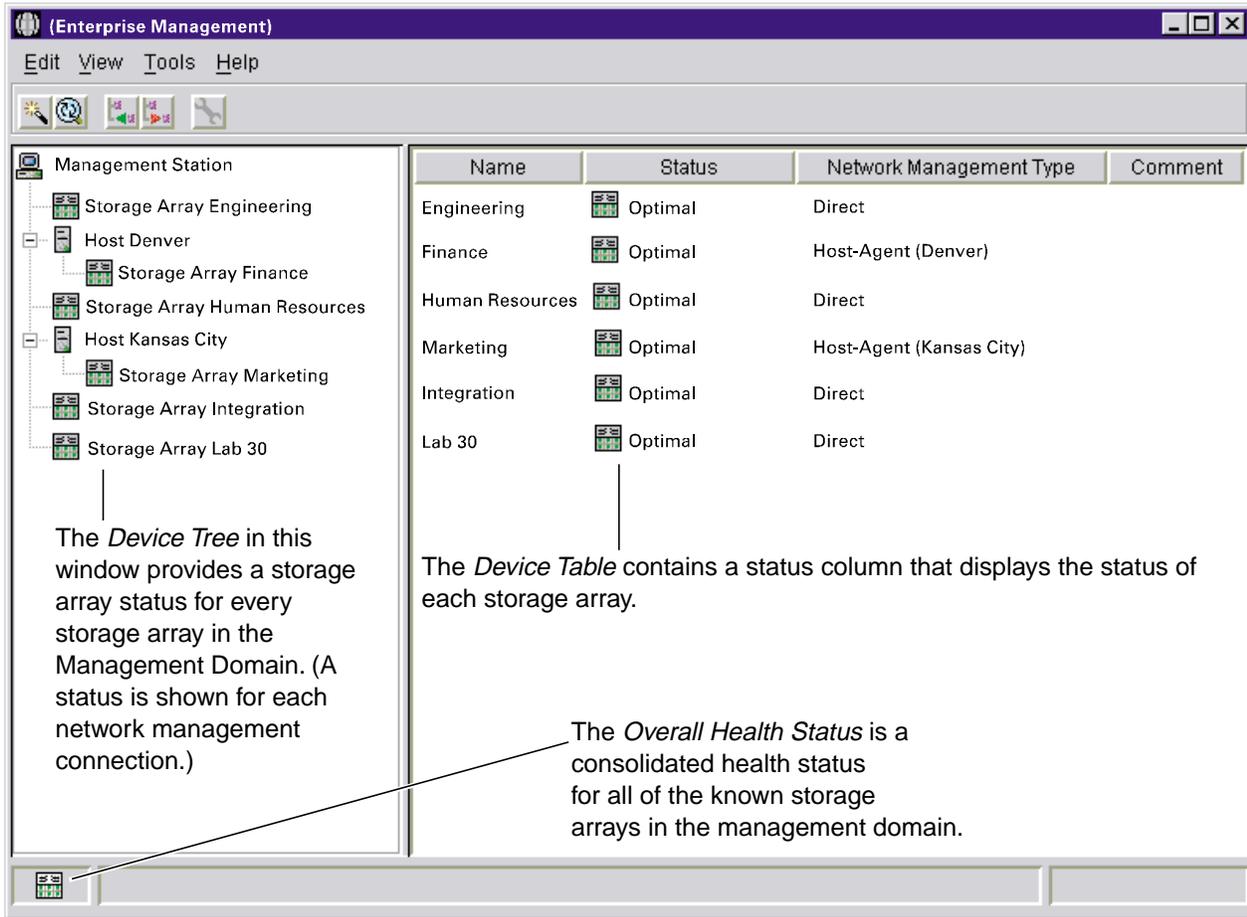
Note: To receive critical alerts, the Enterprise Management window must be running (it can be minimized).

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 43).
- Monitor storage array performance and tune performance, if necessary, (as explained in “Tuning Storage Arrays” on page 48).

Monitoring the Health of Storage Arrays

Figure 2-7 shows the display of storage array status icons in the Enterprise Management window.



33599

Figure 2-7 Monitoring the Health of Storage Arrays

Table 2-5 provides information about the coarse-level storage array status icon that displays 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).
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.

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

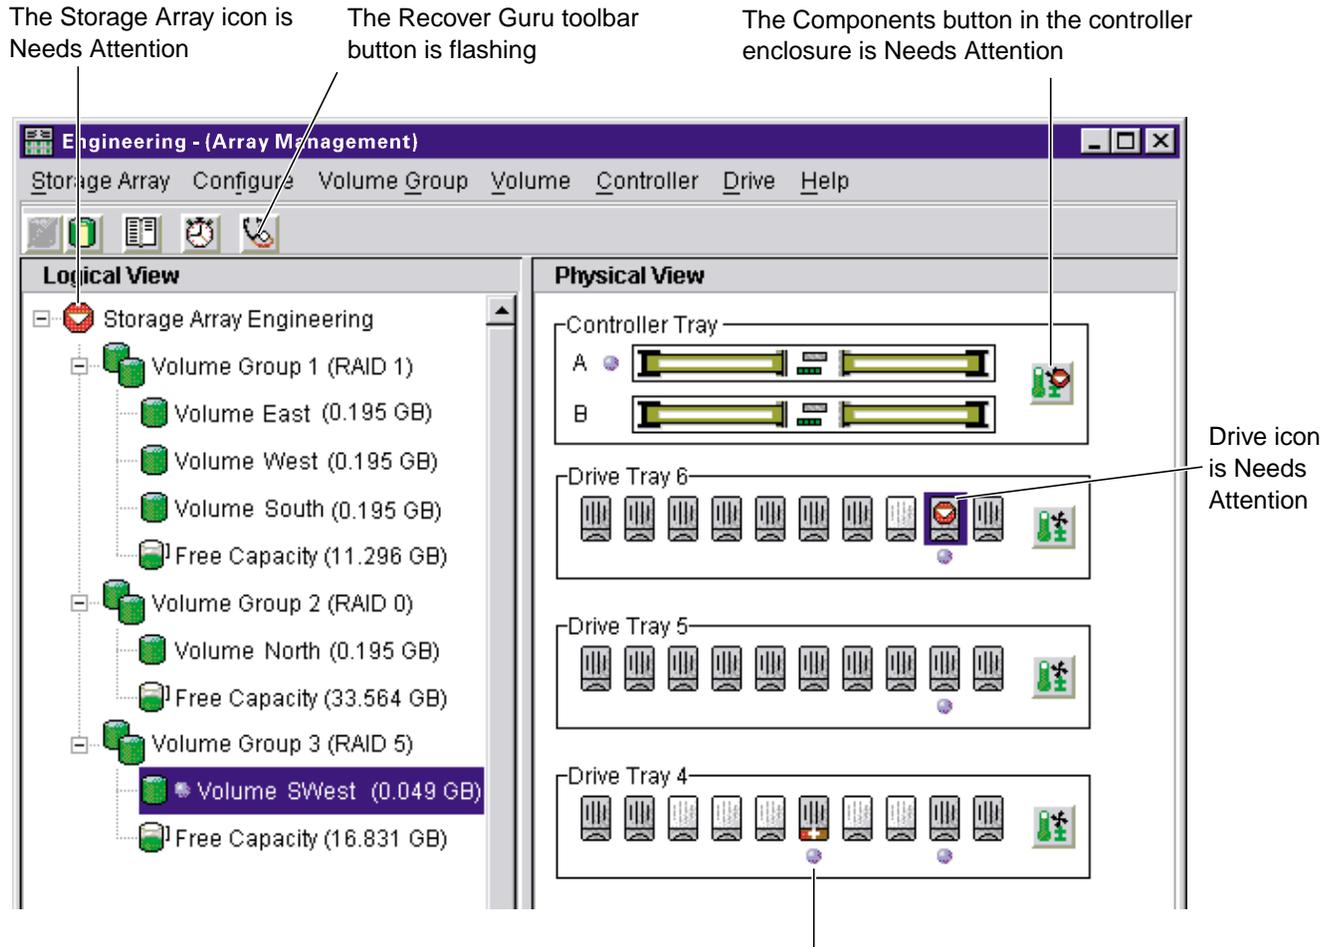
- A Needs Attention storage array icon is displayed in:

- 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 and Physical View.
- You receive critical SNMP or email notifications.

Important: The Enterprise Management window must be running (it can be minimized) for critical alert notifications to be sent.

- Hardware fault lights are displayed.

In Figure 2-8, the Array Management window for storage array Engineering shows indications of a failure in both the Logical View and the Physical View.



The hot spare drive is In Use (taking over for a failed drive). This means that the data from the failed drive has been reconstructed to the hot spare drive. Because the hot spare drive has taken over for the failed drive, the associated volume (Volume SWest) in Volume Group 3 remains Optimal.

Figure 2-8 Failure Notification in the Array Management Window

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

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. Once 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.

Recovering from Storage Array Problems

The **Recovery Guru** is a component of the Array Management window in the SM7client package. The Recovery Guru diagnoses storage array problems and recommends recovery procedures to fix the problems. To display Recovery Guru, select the **Recovery Guru** toolbar button in the Array Management window, called out in Figure 2-9, or select the Storage Array >> Recovery Guru menu option.

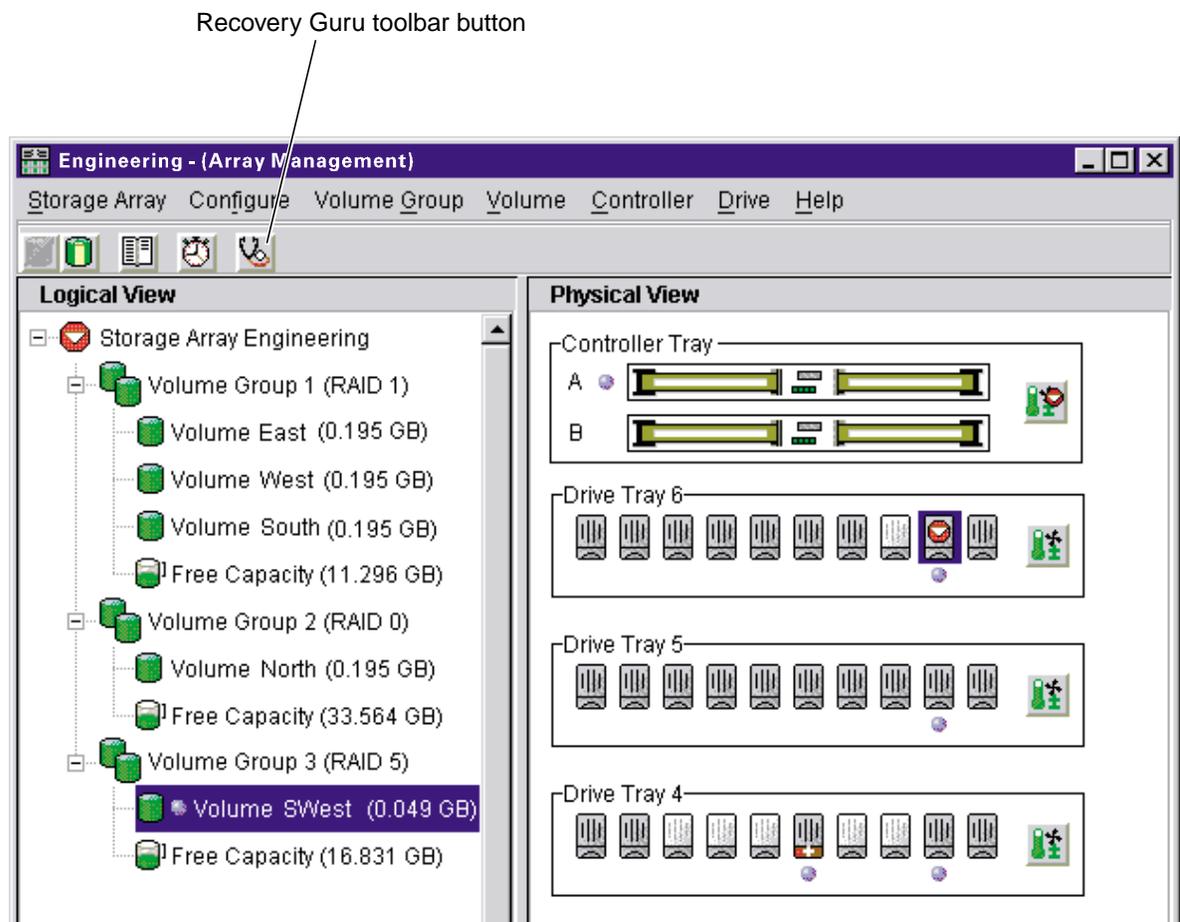


Figure 2-9 Recovery Guru Toolbar Button

The Recovery Guru window is shown in Figure 2-10. The Summary area shows that there are two different failures in this storage array. First, in volume SWest, there is a hot spare drive taking over for a failed drive in enclosure 6, slot 9. Second, there is a failed battery canister in the controller enclosure.

If you select a failure from the list in the Summary area, the appropriate Details and a recovery procedure are shown. Follow the steps in the Recovery Procedure window to recover from each individual failure.

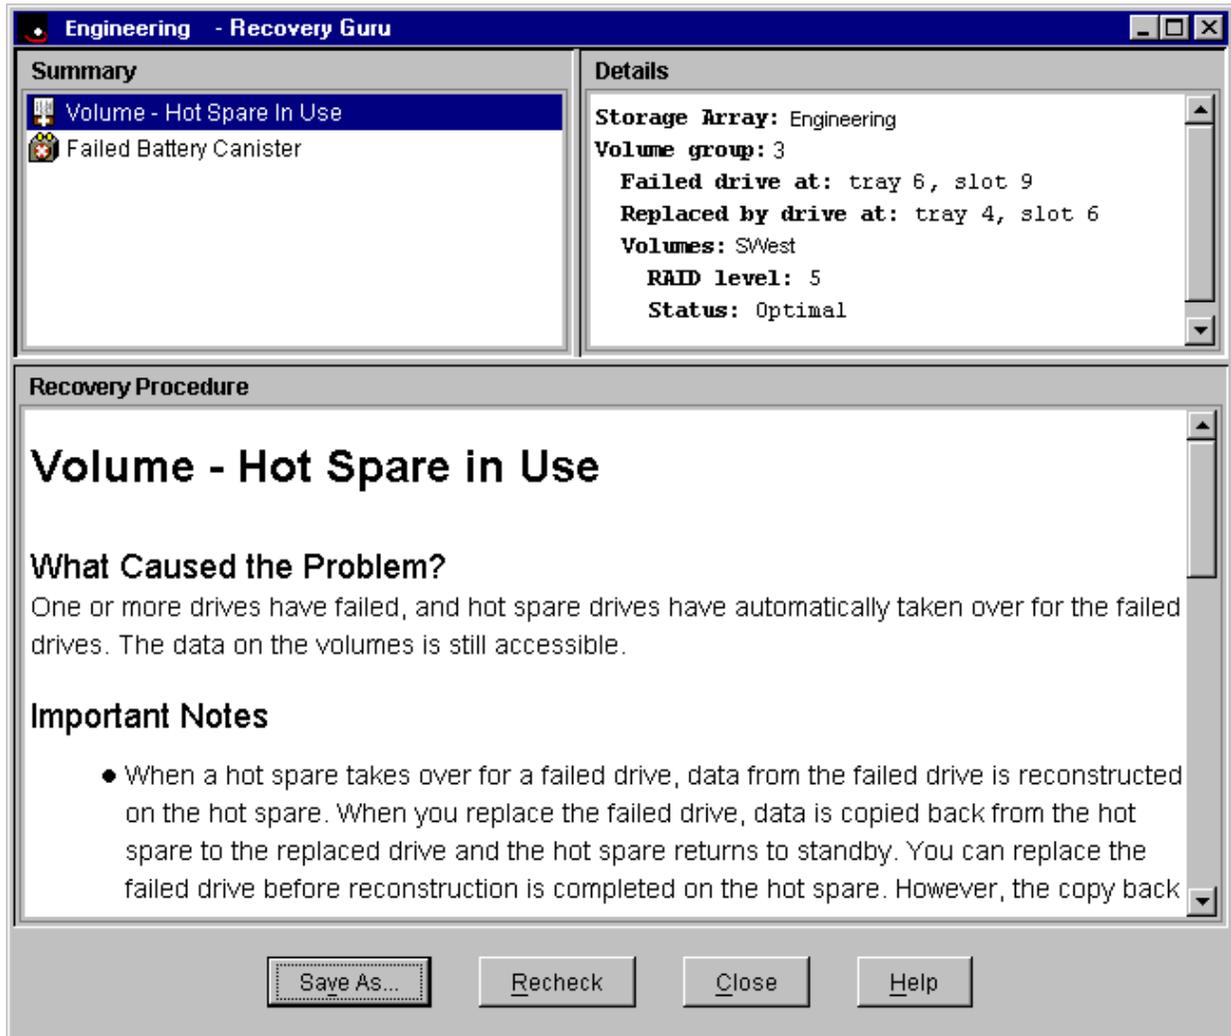


Figure 2-10 Recovery Guru Window

As you follow the recovery procedure to replace the failed drive, the associated volume (SWest) changes status to Operation in Progress and the replaced drive status changes to Replaced Drive. 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 2-11.

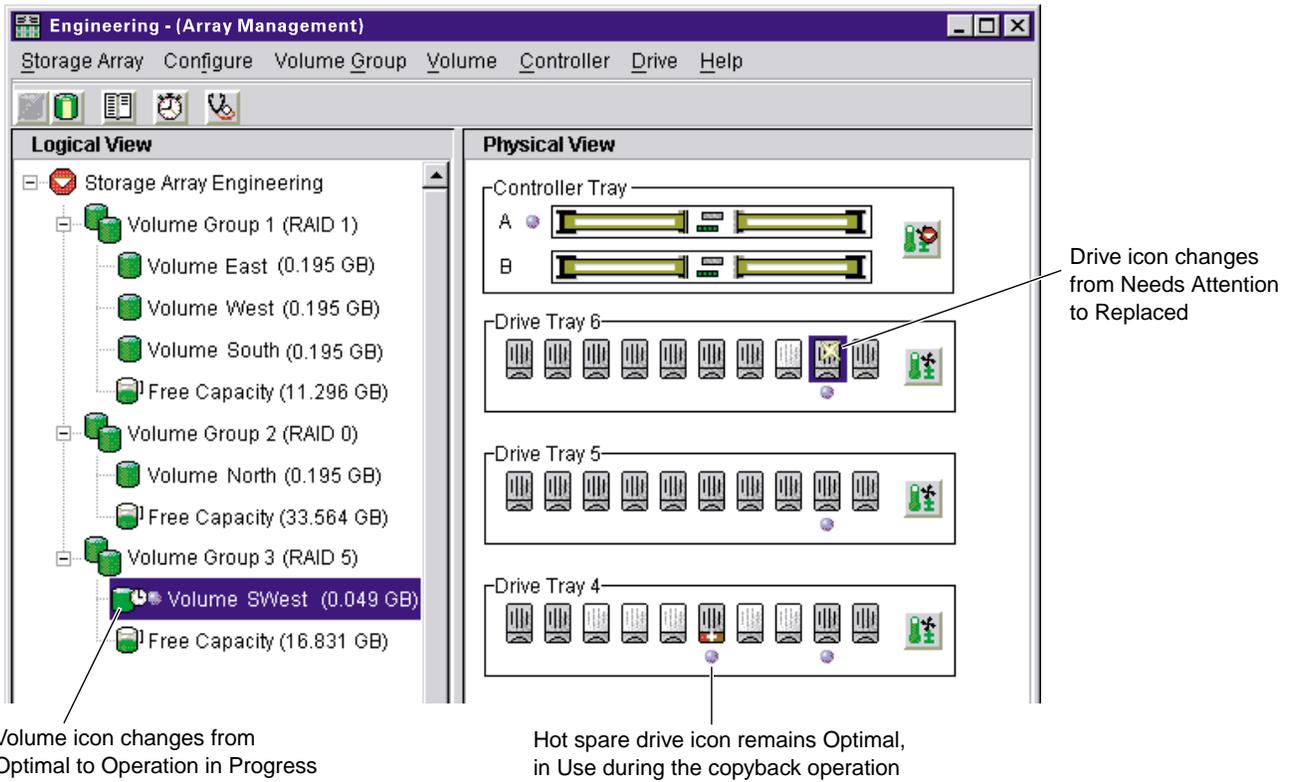


Figure 2-11 Status Changes During Example Recovery Operation

When the copyback operation is finished, the status icons change to reflect the optimal status of the components, as shown in Figure 2-12.

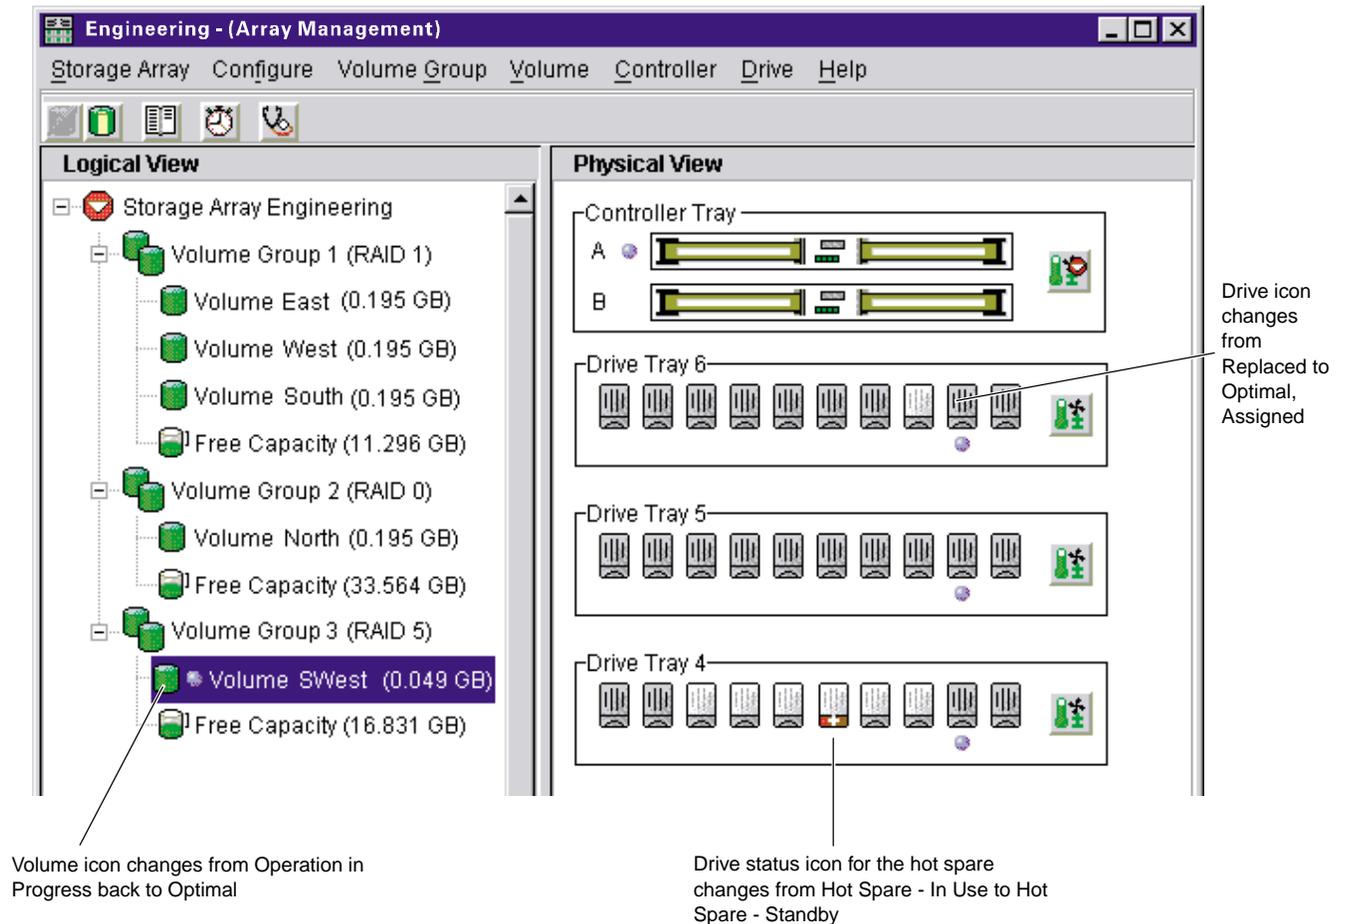


Figure 2-12 Status Changes When an Example Recovery Operation is Completed

After you fix the failed battery canister in the controller enclosure, the following changes occur:

- The **Components** button in the controller enclosure returns to Optimal (Physical View).
- 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.

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 cache hit percentage is the percentage of reads that are serviced with data from the cache rather than requiring a read from disk. - A transfer rate is the amount of data in kilobytes per second that moves through the I/O data path (also called throughput). - An I/O request rate 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 21 and the Array Management window Help.</p>

Table 2-6 (continued) Tuning Options in Array Management Window

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 segment 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 volume group data stripe 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 (continued) Tuning Options in Array Management Window

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>Write caching. 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>Cache read-ahead. 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>Volume modification priority. 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 (continued) Tuning Options in Array Management Window

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 (continued) Task Reference: Storage Array Configuration

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 (continued) Task Reference: Recovering From Problems

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 (continued) Task Reference: Miscellaneous System Administration

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

Index

A

access volume, 7, 8
Add Device option, 12
agent software
 see host-agent software
alert destinations, 37
alert notifications, 10
Array Management Window
 described, 14
 failure notification in, 40
 functions, 14
 introduction to, 9
 monitoring Storage Arrays with, 37
Array Management window Help, xv
Automatic Discovery option, 12

B

battery age clock, 22
BOOTP or BOOTP-Compliant DHCP Server, 2

C

cache block size, 49
cache flush
 described, 22
 performance impacts, 22
 settings, 22

 start flush percentage, 49
 start percentage, 22
 stop flush percentage, 49
 stop percentage, 22
cache hit percentage
 defined, 49
 impact from cache flush settings, 22
cache read-ahead
 choosing a multiplier, 51
 defined, 51
cache settings, 54
channel protection, 23
client software
 where installed, 8
 windows, 9
Components button, 15, 47
configuration
 automatic, 54
 resetting a, 54
configuration file
 described, 10
 recovering a corrupted, 55
Contacting Device status, 39
controller
 assigns default LUNs, 29
 defined, 3
 firmware, 28
 listing components associated with a, 56
 ownership, 54
 placing in active or passive mode, 56
 placing online or offline, 56

- viewing properties, 57
- controller module
 - see controller tray
- controller tray
 - batteries, 22
 - defined, 3
- copyback, 22
- critical alerts, 2
- critical event notification, 24

D

- data protection strategies, 19
- Default Host Group
 - defined, 28
 - when defined host groups or host are part of, 33
- default LUN, 28
- default volume-to-LUN mapping
 - and volumes, 35
 - defined, 30
- device
 - automatic discovery of, 52
 - correcting a partially managed, 55
 - manually adding, 52
 - removing, 52
- Device Table, 38
- Device Tree, 38
- directly managed storage array, 7
- drive
 - defined, 3
 - failing, 55
 - initializing, 56
 - listing components associated with, 56
 - locating, 56
 - reconstructing, 56
 - reviving, 56
 - viewing properties, 57
- drive group

- see volume group
- drive module
 - see drive tray
- drive tray
 - defined, 3
 - Hardware Installation and User Guides, xiv

E

- email destinations, 55
- emwdata.bin file, 10
- Enterprise Management Window
 - figure, 11
 - functions, 10
 - introduced, 9
 - monitoring storage arrays with, 37
 - status icons displayed in, 38
- Enterprise Management window Help, xv
- Event Log, 55
- event notification
 - configuring alert destinations, 24
- expected usage, of a volume, 25

F

- failure notification
 - in the Array Management Window, 40
 - types of, 39
- firmware
 - assigns default LUNs, 29
 - downloading, 53
- Fixing status, 39
- free capacity, 4
 - adding to a volume group, 25, 53
 - consolidating on a volume group, 53
 - creating volumes from, 53
 - creating volumes on, 25

defined, 4
increasing storage array, 53
Free Capacity node, 4

H

hardware components, 1
Help
 task reference, 52
 using, xv
host
 defined, 2
 defined as a topological element, 28
 view of partitioned storage, 33
host group
 defined, 28
 view of partitioned storage, 33
host port
 defined, 29
 discovery of, 29
host-agent managed Storage Array
 described, 7
 discovering newly attached, 52
host-agent software
 described, 8
 package contents, 8
 where installed, 8
hot spare drive
 assigning, 53
 configuring, 22
 defined, 22
 unassigning, 53

I

indicator lights, 56
Installation and Support Guide, xiv

I/O access pattern, 50
I/O data connection
 see I/O data path
I/O data path
 described, 5
I/O request rate
 defined, 49
 impact from cache flush settings, 22
 modification priority effect on, 52

L

logical unit number
 see LUN
Logical View, 39
 described, 15
 status icon changes, 45, 46
LUN
 address space, 29, 35
 defined, 29
 using in a volume-to-LUN mapping, 35
 see volume

M

management domain
 adding host-agent managed storage arrays to, 7
 configuration file, 10
 defined, 5
 global e-mail alert settings, 55
 populating, 12
 use of Automatic Discovery, 12
management station
 see storage management station, 1
Mappings View
 described, 27
 volume-to-LUN mapping in, 36
media scan

changing settings, 54
defined, 23
medical imaging applications, 21
modification operations, 51
modification priority
 changing, 54
 described, 51
 effect on I/O request rate, 52
 setting, 52
multimedia applications, 21
multi-user environments, 21

N

Needs Attention status, 39
network management connection
 combinations of, 7
 defined, 6
 host-agent, 8
 types of, 6
Network Management Station
 defined, 2
 receipt of failure notifications, 42
NVSRAM, downloading, 53

O

online Help
 see Help
Optimal status, 39
Overall Health Status, 38

P

password protection, 24, 56
Performance Monitor, 49, 54

Physical View
 changing tray order in, 53
 described, 15
 status icon changes, 45, 46
power outage, 22

R

RAID 0
 described, 20
 drive failure consequences, 20
RAID 1
 described, 20
 drive failure consequences, 20
 performance, 50
 redundancy, 20
RAID 3
 described, 21
 drive failure consequences, 21
 redundancy, 20
RAID 5
 described, 21
 drive failure consequences, 21
 performance, 50
 redundancy, 20
RAID level
 and channel protection, 23
 changing, 50, 54
 described, 19
 use on a volume group, 20
 using an appropriate, 19
RAID Module
 see Storage Array
read percentage, 50
reconstruction, 22
Recovery Guru
 described, 43
 identifying when to use, 55

Recovery Procedure, 44
Summary Area, 44

redundancy
checking, 55
data, 19
of Fibre Channel Arbitrated Loops, 23
of Storage Array hardware, 19

remaining capacity
see free capacity

result, 36

S

Script Editor, 10, 53

segment, 50

segment size
changing, 50, 54
selecting an appropriate, 50

SM7agent
see host-agent software

SM7client
description, 8
windows in, 9

SM7devices utility, 8

SNMP destinations
specifying, 55
where stored, 10

SNMP trap message
interpreting, 55
use for critical alerts, 2

software components, 8

specific volume-to-LUN mapping, 35

start percentage, cache flush, 22

status icon
changes during a recovery procedure, 46
changes during recovery procedure, 45
displayed in the Enterprise Management window,
37

stop percentage, cache flush, 22

Storage Array
comment, 52
configuring, 19
connectivity, 5
defined, 2
directly managed, 7
host-agent managed, 7
indications of failure, 39
locating, 56
logical components, 4
management activities, 5
physical components, 3
premium feature, 26
renaming, 56
synchronizing clocks, 57
tuning options available, 48
viewing properties, 57

Storage Array Management, 1

storage management software
where installed, 1

storage management station
defined, 2
described, 1

Storage Partition
configuring, 32
creating, 24, 53
defined, 4
Definitions window, 26
major steps to creating, 32
Topology, 28, 33

T

task reference, 52

terminology
functional equivalents, xiii
text formatting conventions, xiii

topological elements

- defined, 27
- display of, 33
- when to define, 28

Topology View

- described, 27
- host group with specific mappings, 36

transfer rate, 49

U

- unassigned capacity
 - see unconfigured capacity
- unconfigured capacity
 - adding to a storage array, 53
 - creating volumes from, 53
 - creating volumes on, 25
 - defined, 4
- Unconfigured Capacity Node, 4
- Understanding guide
 - how to use, xiii
 - organization, xiii
 - prerequisite knowledge, xiii
- Unresponsive status
 - defined, 39
 - recovery, 55
- UPS, 22

V

- volume
 - creating, 24
 - defined, 4
 - initializing, 56
 - renaming, 56
 - sharing access to a, 35
 - viewing properties, 57
- volume group
 - capacity of, 20

- defined, 4
- deleting, 53
- initializing, 56
- locating, 56
- moving to another controller, 50
- redistributing, 56
- volume group data stripe, 50
- volume modification priority
 - see modification priority
- volume-to-LUN mapping
 - default, 30, 32, 35
 - defined, 34
 - None (map later) setting, 32
 - shared, 35
 - specific, 30, 35
- volume-to-LUN mapping preference, 25

W

- worldwide name of host port, 29
- write cache mirroring
 - described, 21, 22
 - how to enable, 21, 22
- write caching
 - and data loss, 21
 - and performance, 21
 - described, 51
- write caching without batteries, 22