

BladeStoreTM Storage Manager

Concepts Guide for Version 8.36

P/N 96090, First Edition



Proprietary Rights Notice

This document contains proprietary information of LSI Logic and Storage Technology Corporation. The information contained herein is not to be used by or disclosed to third parties without the express written permission of an officer of LSI Logic Corporation or Storage Technology Corporation. Any product(s) described herein is/are a licensed product of LSI Logic Corporation and Storage Technology Corporation.

Document Description

Document 12272-E1, First Edition. September 2002

This document describes version 8.3x of BladeStore™ Storage Manager, and will remain the official reference source for all revisions/releases of this product until rescinded by an update.

Intended Readers

This book is intended for end users, system operators, system administrators, and service technicians who are responsible for installing hardware. Readers should have knowledge of hardware installation and operation and understand Redundant Array of Independent Disks (RAID), Small Computer System Interface (SCSI), and Fibre Channel (FC).

Disclaimer

It is the policy of LSI Logic and Storage Technology Corporation to improve products as new technology, components, software, and firmware become available. We reserve the right to make changes to any products herein at any time without notice. All features, functions, and operations described herein may not be marketed in all parts of the world. In some instances, photographs and figures are of equipment prototypes. Therefore, before using this document, consult your sales representative or account team for information that is applicable and current. WE DO NOT ASSUME ANY RESPONSIBILITY OR LIABILITY FOR THE USE OF ANY PRODUCT(S) DESCRIBED HEREIN EXCEPT AS EXPRESSLY AGREED TO IN WRITING BY LSI LOGIC.

License Restriction

The purchase or use of an LSI Logic/StorageTek solution does not convey a license under any patent, copyright, trademark, or other intellectual property right of LSI Logic, StorageTek, or its third parties.

Copyright Notice

© 2002 LSI Logic Corporation. All rights reserved.

© 2002 Storage Technology Corporation. All rights reserved.

Trademark Acknowledgments

LSI Logic, the LSI Logic logo, StorageTek, the StorageTek logo, ContinuStor, and SANtricity are all trademarks or registered trademarks of LSI Logic Corporation or Storage Technology Corporation. All other brand and product names may be trademarks of their respective companies.

Regulatory Compliance Statements

FCC Radio Frequency Interference Statement

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the Federal Communications Commission (FCC) Rules. These limits are designed to provide reasonable protection against harmful interference in a commercial installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

LSI Logic Corporation and Storage Technology Corporation are not responsible for any radio or television interference caused by unauthorized modification of this equipment or the substitution or attachment of connecting cables and equipment other than those specified by LSI Logic Corporation and Storage Technology Corporation. It is the user's responsibility to correct interference caused by such unauthorized modification, substitution, or attachment.

Laser Products Statement

This equipment has been tested and found to comply with regulations for Class 1 laser product pursuant to 21 CFR, Section 1040-10. For outside the USA, this equipment has been tested and found compliant with Class 1 laser product requirements contained in European Normalization standard EN 60825-1 1994+A11. Class 1 levels of laser radiation are not considered to be hazardous and are considered safe based upon current medical knowledge. This class includes all lasers or laser systems which cannot emit levels of optical radiation above the exposure limits for the eye under any exposure conditions inherent in the design of the laser product.

LSI Logic Corporation and Storage Technology Corporation are not responsible for any damage or injury caused by unauthorized modification of this equipment or the substitution or attachment of connecting cables and equipment other than those specified by LSI Logic Corporation and Storage Technology Corporation. It is the user's responsibility to correct interference caused by such unauthorized modification, substitution, or attachment.

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classé A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

この装置は、第一種情報装置（商工業地域において使用されるべき情報装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会（VCCI）基準に適合しております。

従って、住宅地域またはその隣接した地域で使用すると、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

警告使用者： 這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Revision Record

Edition or Revision	Date	Affected Pages or Remarks
First Edition	October 2002	New Book.

Part Number: 96090

Contents

Chapter 1: ABOUT THIS BOOK

Chapter 2: INTRODUCTION TO THE SOFTWARE

Client Software Windows	2-2
The Enterprise Management Window.....	2-4
Device Tree	2-6
Device Table	2-6
Enterprise Management Window Menus.....	2-7
Enterprise Management Window Toolbar	2-8
The Array Management Window	2-9
Array Management Window Tabs	2-10
Array Management Window Menus	2-12
Array Management Window Toolbar	2-13

Chapter 3: PROTECTING YOUR DATA

Software Redundancy	3-2
Remote Volume Mirroring.....	3-2
Snapshot Volumes.....	3-2
RAID Levels and Data Redundancy	3-3
Hardware Redundancy	3-5
Controller Cache Memory	3-5
Hot Spare Drives	3-6
Background Media Scan	3-6
Channel Protection.....	3-7
I/O Data Path Protection	3-8
Multi-Path Driver With AVT Enabled.....	3-8
Multi-Path Driver with AVT Disabled	3-8
Password Protection	3-9

Chapter 4: CONFIGURING DISK SUBSYSTEMS

Volumes and Volume Groups	4-2
Volume Groups	4-3
Standard Volumes	4-4
Specifying Volume Parameters from Free Capacity	4-4
Specifying Volume Parameters from Unconfigured Capacity	4-5
Snapshot Volumes.....	4-6
Disabling or Deleting a Snapshot Volume	4-7
Re-creating a Snapshot Volume	4-7
Remote Volume Mirroring.....	4-8
Mirror Relationships	4-8
Mirror Repository Volumes	4-8
Data Replication	4-9
Register the Volume with the Operating System.....	4-9
Dynamic Volume Expansion (DVE)	4-10
Mappings View	4-11
SANshare Storage Partitioning	4-14
SANshare Storage Partitioning Example	4-14
Heterogeneous Hosts.....	4-17
Heterogeneous Hosts Example	4-17

Chapter 5: MAINTAINING AND MONITORING DISK SUBSYSTEMS

Disk Subsystem Health	5-2
Disk Subsystem Status Icons	5-3
Event Monitor.....	5-4
Alert Notifications.....	5-6
Customer Support Alert Notifications	5-7
Problem Notification	5-8
Disk Subsystem Problem Recovery	5-10
Recovery Guru Example	5-11

List of Figures

Chapter 2: INTRODUCTION TO THE SOFTWARE

Figure 2-1. Enterprise Management Window	2-2
Figure 2-2. Array Management Window	2-3
Figure 2-3. Enterprise Management Window	2-5
Figure 2-4. Device Tree of a Management Domain Example	2-6
Figure 2-5. Array Management Window	2-10

Chapter 4: CONFIGURING DISK SUBSYSTEMS

Figure 4-1. Unconfigured and Free Capacity Nodes	4-3
Figure 4-2. Mappings View Window	4-11
Figure 4-3. SANshare Storage Partitioning Example	4-15
Figure 4-4. Host Port Definitions Dialog	4-17
Figure 4-5. Heterogeneous Hosts Example	4-18

Chapter 5: MAINTAINING AND MONITORING DISK SUBSYSTEMS

Figure 5-1. Monitoring Disk Subsystem Health Using the Enterprise Management Window	5-2
Figure 5-2. Event Monitor Configuration	5-4
Figure 5-3. Event Monitor Example	5-5
Figure 5-4. Problem Notification in the Array Management Window	5-9
Figure 5-5. Displaying the Recovery Guru Window	5-10
Figure 5-6. Recovery Guru Window Example	5-11
Figure 5-7. Status Icon Changes During an Example Recovery Operation	5-12
Figure 5-8. Status Changes When The Example Recovery Operation is Completed	5-13

List of Tables

Chapter 2: INTRODUCTION TO THE SOFTWARE

Table 2-1. Enterprise Management Window Menus	2-7
Table 2-2. Enterprise Management Window Toolbar Buttons	2-8
Table 2-3. Array Management Window Tabs	2-11
Table 2-4. Array Management Window Menus	2-12
Table 2-5. Array Management Window Toolbar Buttons	2-13

Chapter 3: PROTECTING YOUR DATA

Table 3-1. RAID Level Configurations	3-3
--	-----

Chapter 4: CONFIGURING DISK SUBSYSTEMS

Table 4-1. Mappings View Tab	4-11
Table 4-2. Volume-to-LUN Terminology	4-12

Chapter 5: MAINTAINING AND MONITORING DISK SUBSYSTEMS

Table 5-1. Disk Subsystem Status Icon Quick Reference	5-3
---	-----

About This Book

This book provides the conceptual framework necessary to understand the main features and functions of the storage management software, including a description of the basic software layout, an introduction to software concepts, information about disk subsystem configuration options, descriptions of functions of the storage management software, and information about monitoring disk subsystems in a management domain.

The Array Management Window and the Enterprise Management Window, outlined in Chapter 2, each contain online help systems that provide context-sensitive procedural and conceptual information. Refer to online help for additional information about the features presented in the following chapters.

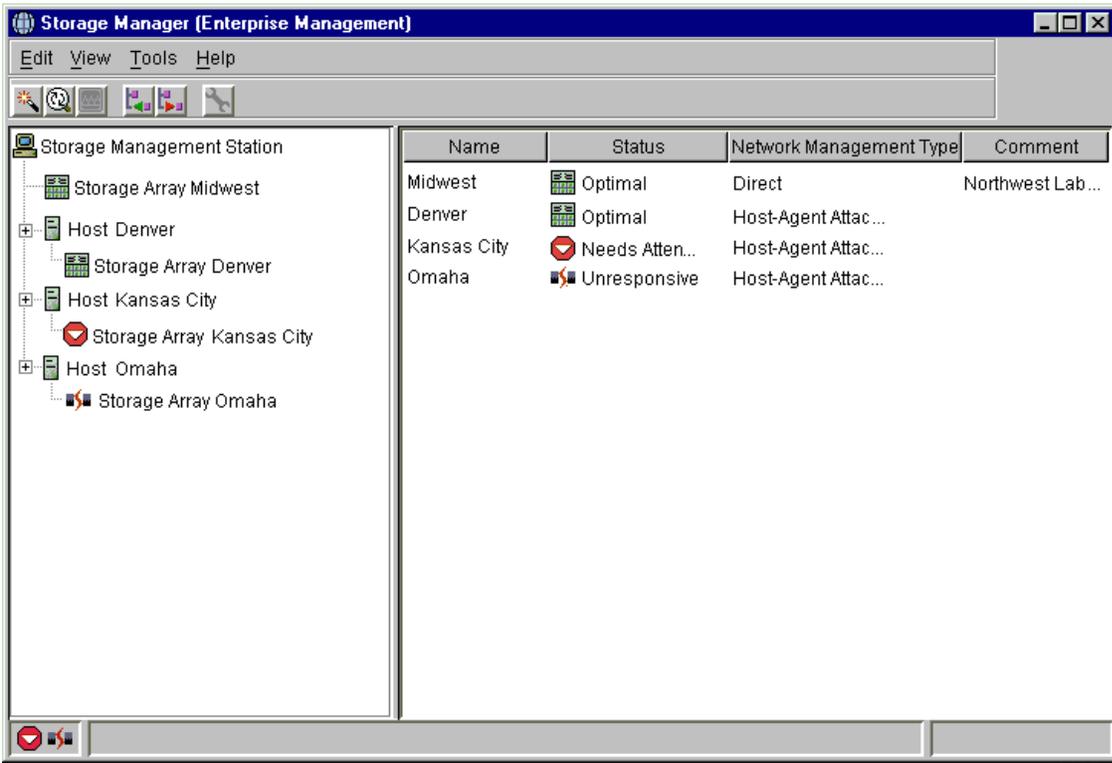
Introduction to the Software

This chapter provides introductions to the Enterprise Management Window and the Array Management Window and describes the basic software layout. This chapter also provides an overview of the tools and functions that are performed from each management window, outlines the menus and toolbar buttons available in each management window, and describes the various views and tabs displayed in the Array Management Window.

For additional conceptual information and detailed procedures for the options described in this section, refer to Storage Manager Concepts, “Introducing the Enterprise Management Window” in the Enterprise Management Window online help, and Learn About Storage Management Software, “Using the Array Management Window” in the Array Management Window online help.

Client Software Windows

The client software has two main windows: the Enterprise Management Window (Figure 2-1) and the Array Management Window (Figure 2-2).



63157

Figure 2-1 Enterprise Management Window

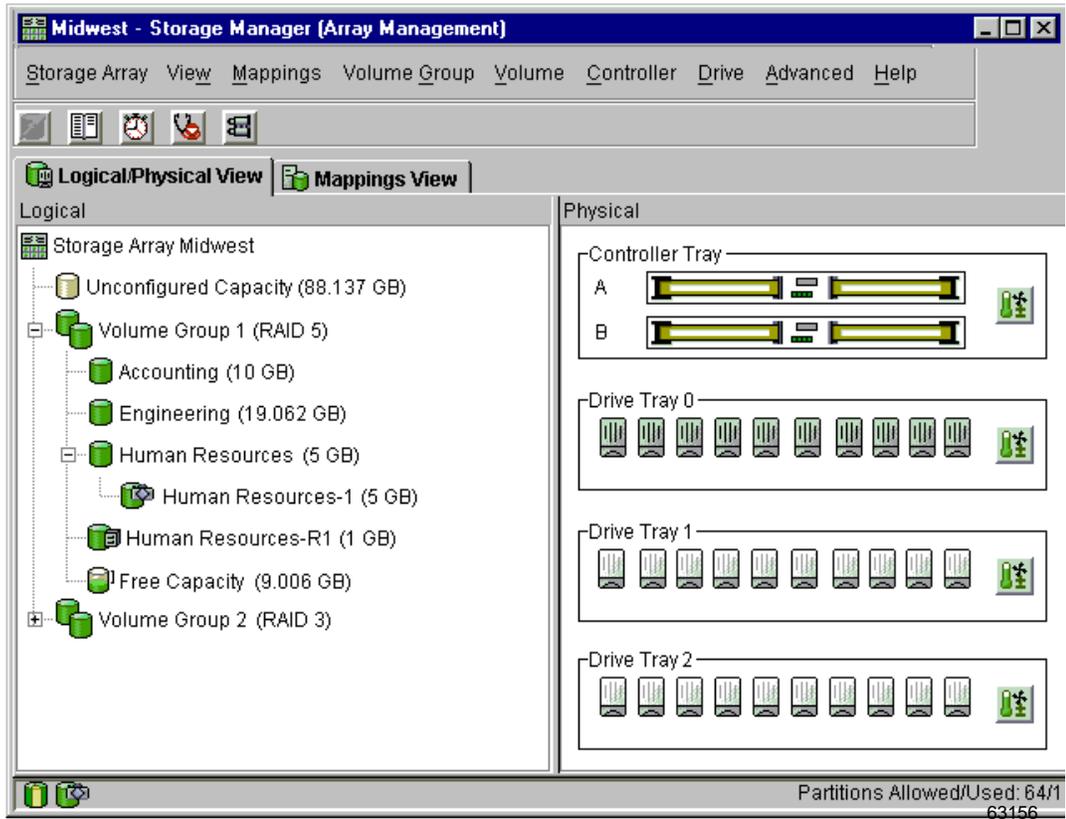


Figure 2-2 Array Management Window

The Enterprise Management Window

The Enterprise Management Window ([Figure 2-3](#)) is the first window to appear when you start the software. It is used to:

- Add and detect the disk subsystems you want to manage.
- Provide a comprehensive view of all disk subsystems in the management domain.
- Perform batch disk subsystem management tasks using the Script Editor.
- Configure destinations, such as e-mail or Simple Network Management Protocol [SNMP] traps, to receive alert notifications for non-optimal disk subsystems.

A local configuration file stores all the information about the management domain, including the disk subsystems you have added and the e-mail or SNMP destinations you have configured. This user-specific configuration file allows each system administrator to create a unique view of the devices in the management domain, without affecting the information shown to other system administrators.

The Enterprise Management Window or the event monitor must be running to receive alert notifications of critical events on disk subsystems.

After disk subsystems are added, the Enterprise Management Window is used primarily to monitor the disk subsystems in the management domain for a Needs Attention status and for alert notification of critical errors affecting the disk subsystems. When you are notified of a non-optimal disk subsystem status in the Enterprise Management Window, starting an Array Management Window for the affected disk subsystem displays detailed information about the disk subsystem condition.

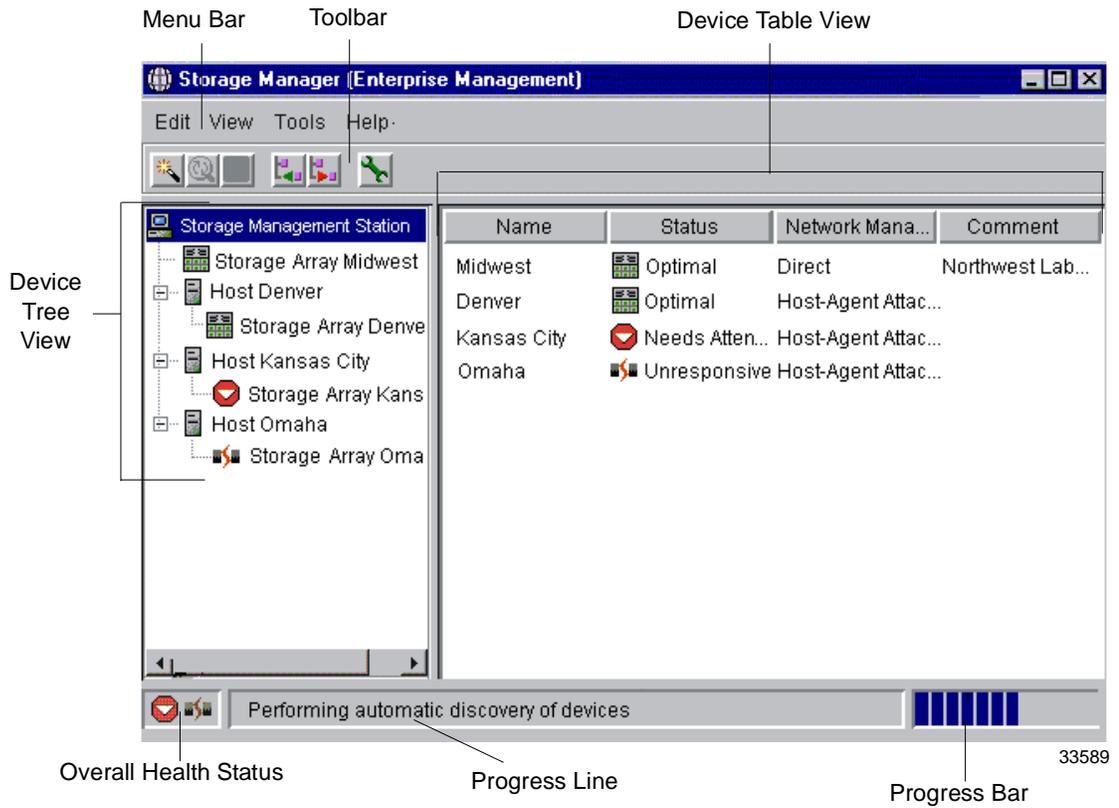


Figure 2-3 Enterprise Management Window

Device Tree

The Enterprise Management Window Device Tree displays the management domain (Figure 2-4). The Device Tree provides a hierarchical view of the nodes in the management domain. The storage management station node is the root node and sends the storage management commands. The Device Tree displays the configured management domain.

When disk subsystems are added to the management domain, they are shown in the Device Tree as child nodes of the storage management station node. A disk subsystem can be managed through an Ethernet connection on each controller in the disk subsystem (direct managed) or through a host interface connection to a host with the host-agent installed (host-agent managed).

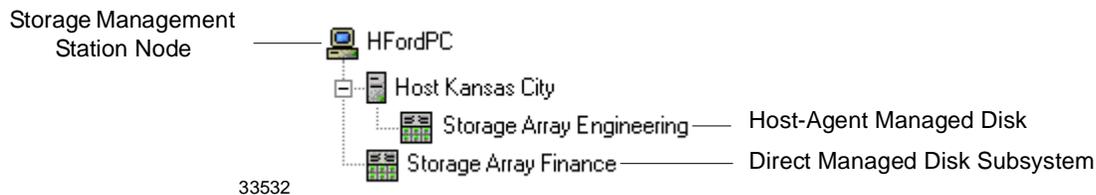


Figure 2-4 Device Tree of a Management Domain Example

There are two ways to populate a management domain:

- **Automatic Discovery** – detects direct managed and host-agent managed disk subsystems on the local subnetwork and adds them to the management domain automatically. The Enterprise Management Window detects host-agent managed disk subsystems by locating any hosts that provide network management connections to the disk subsystems. The hosts appear in the Device Tree with their associated disk subsystems.
- **Add Device** – provides a way to manually add direct managed and host-agent managed disk subsystems to the management domain. For a direct managed disk subsystem, a host name or IP address must be entered for each controller in the disk subsystem. Typically, there are two controllers in a single disk subsystem. For a host-agent managed disk subsystem, a name or an IP address must be entered for the host that is attached to the disk subsystem.

The first time you detect or add disk subsystems to a management domain, they are shown as <unnamed> in the Device Tree and Device Table unless they have been named by another storage management station.

Device Table

The Device Table lists the name, status, management type (direct network attached for direct managed disk subsystems, or host-agent attached for host-agent managed disk subsystems), and comments entered for disk subsystems (Figure 2-3 on page 2-5).

Enterprise Management Window Menus

The Enterprise Management Window menus on the menu bar are described in [Table 2-1](#).

Table 2-1 Enterprise Management Window Menus

Menu	Description
Edit	Contains options to add devices to or remove devices from the management domain, to configure alert destinations, or to add a comment to the Device Table about a disk subsystem.
View	Provides options to sort the entries in the Device Table by name, status, management type, or comment. Another option shows partially managed devices—a condition occurring when only one controller of a controller pair is defined or can be reached when the disk subsystem is added or detected. In this state, many management operations that require access to both controllers are not available.
Tools	Displays options for automatically detecting devices on the same subnetwork or for rescanning to find disk subsystems newly attached to a host. Other items on this menu include options to update the event monitor, open an Array Management Window to manage a selected disk subsystem, open the Script Editor to perform batch management tasks, or load a saved configuration file.
Help	Presents options to display the Enterprise Management Window Help system and to view the software version and copyright information.

Enterprise Management Window Toolbar

The Enterprise Management Window toolbar buttons are described in [Table 2-2](#).

Table 2-2 Enterprise Management Window Toolbar Buttons

Toolbar Button	Description
Automatically detect new devices 	Activates the Automatic Discovery option that detects hosts and disk subsystems on the local subnetwork and adds them to the management domain.
Rescan selected host for new devices 	Rescans the highlighted host for any newly attached disk subsystems. Before using this option, the new disk subsystems must be physically attached to the host and the host-agent software residing on the host must be restarted. Note This option is available only when you select a host in the Device Tree.
Synchronize Event Monitor 	Synchronizes the event monitor with any changes made in the Enterprise Management Window, such as adding or removing devices or adding alerts. Note This option is available only if the configurations of the Enterprise Management Window and the event monitor are not synchronized.
Add host/device 	Opens the Add Device dialog for manually adding hosts or disk subsystem controllers to the management domain.
Remove host/device 	Removes the selected disk subsystem or the selected host and its attached disk subsystems from the management domain. Note This option is available only when a disk subsystem in the Device Tree or Device Table is selected, or a host in the Device Tree is selected.
Launch an Array Management Window 	Starts an Array Management Window for the selected disk subsystem. Note This option is available only when a disk subsystem in the Device Tree or Device Table is selected.

The Array Management Window

The Array Management Window (Figure 2-5 on page 2-10) is used to configure and maintain the logical and physical components of a disk subsystem and to view and define volume-to-LUN mappings.

The Array Management Window is specific to an individual disk subsystem; therefore, you can manage only a single disk subsystem within an Array Management Window. However, you can start other Array Management Windows from the Enterprise Management Window to simultaneously manage multiple disk subsystems.

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

The features of a particular release of firmware will be accessible when an Array Management Window is launched from the Enterprise Management Window to manage a disk subsystem. For example, you manage two disk subsystems using this software; one disk subsystem has firmware version 5.30 and the other has firmware version 4.x. When you launch an Array Management Window for a particular disk subsystem, the correct Array Management Window version is used. The disk subsystem with firmware version 5.30 will use version 8.30 of the storage management software while the disk subsystem with firmware version 4.x will use version 7.x. You can verify the version you are currently using by selecting Help >> About in the Array Management Window.

This bundling of previous versions of the Array Management Window provides the flexibility of upgrading the firmware only on selected disk subsystems instead of having to perform an upgrade on all disk subsystems at once.

Array Management Window Tabs

The Array Management Window has two tabs: Logical/Physical View and Mappings View (Figure 2-5), which are described in Table 2-3. The tabs display information about the logical components (volume and volume groups), physical components (controllers and drives), topological elements (host groups, hosts, host ports), and volume-to-LUN mappings in a disk subsystem.

Also displayed in the Array Management Window are the toolbar, menu bar, components buttons, premium feature status area, and the storage partition status area.

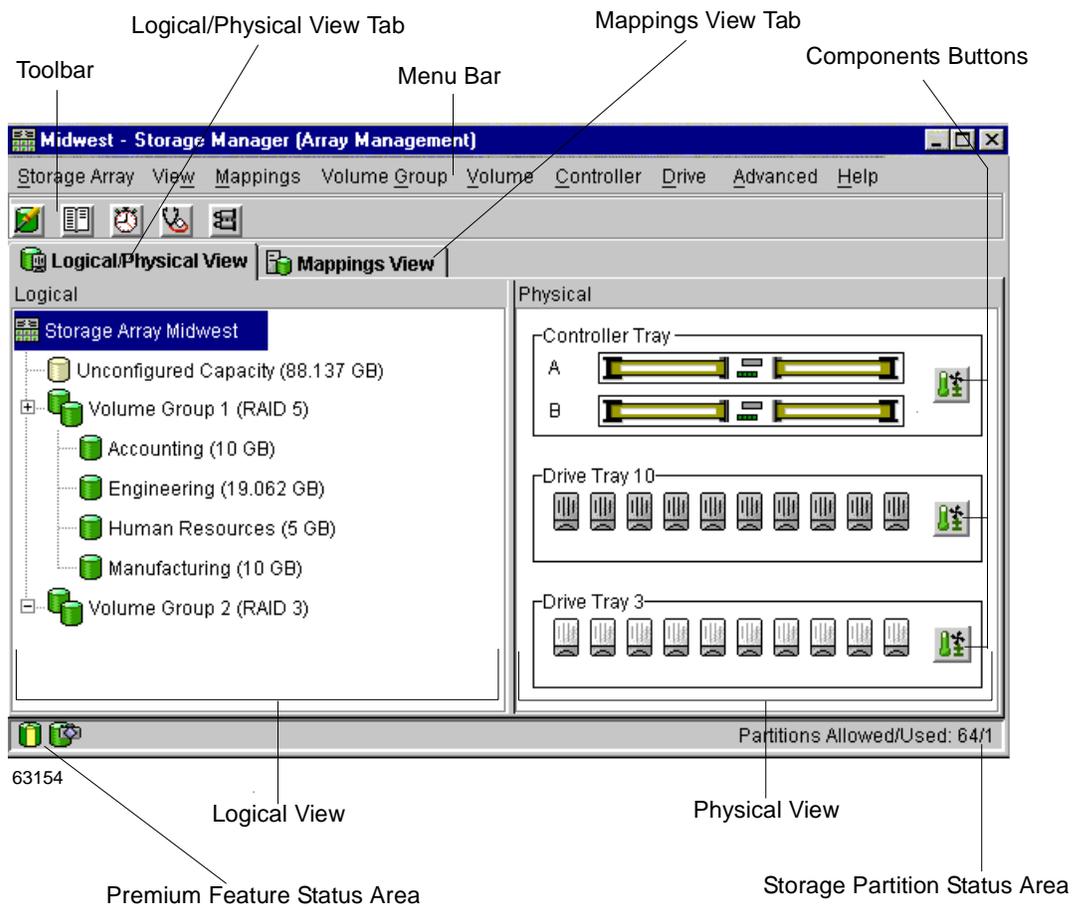


Figure 2-5 Array Management Window

Table 2-3 Array Management Window Tabs

Tabs	Description
Logical/Physical View	<p>The Array Management Window Logical/Physical View contains two panes: the Logical View and the Physical View.</p> <p>The Logical View (left pane of Figure 2-5) provides a tree-structured view of logical nodes. This view shows the organization of disk subsystem capacity into volume groups and volumes.</p> <p>The Physical View (right pane of Figure 2-5) provides a view of the physical devices in a disk subsystem, such as control module and disk array components.</p> <p>Selecting a volume or other entity in the Logical View shows you the associated physical components in the Physical View.</p> <p>There is a Components button in every control module and disk array that, when selected, presents the status of each component and shows the temperature status.</p>
Mappings View	<p>The Mappings View of the Array Management Window contains two panes: the Topology View and the Defined Mappings View. For details, see “Mappings View” on page 4-11.</p> <p>The Topology View provides a tree-structured view of logical nodes related to storage partitions.</p> <p>The Defined Mappings Detail View displays the mappings associated with the selected node in the Topology View.</p>

Array Management Window Menus

The Array Management Window menus are described in [Table 2-4](#). The menus are used to perform storage management operations for a selected disk subsystem or for selected components within a disk subsystem. See [Figure 2-5 on page 2-10](#) for an example of the Array Management Window menus.

Table 2-4 Array Management Window Menus (1 of 2)

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

Table 2-4 Array Management Window Menus (2 of 2)

Menu	Description
Drive	Contains options to perform the following storage management operations on drives: locating a drive, assigning or unassigning a hot spare, failing, reconstructing, reviving or initializing a drive, or viewing drive properties.
Advanced	Presents maintenance options which should only be used under the guidance of technical support.
Help	Provides options to display the contents of the Array Management Window Help, view a reference of all Recovery Guru procedures, and to view the software version and copyright information.

Array Management Window Toolbar

The Array Management Window toolbar buttons are used to create new volumes or volume groups, monitor performance, view events, recover from failures, and locate a particular node. Each toolbar button is described in [Table 2-5](#). See [Figure 2-5 on page 2-10](#) for an example of the Array Management Window toolbar.

Table 2-5 Array Management Window Toolbar Buttons

Toolbar Button	Description
Create new volumes (and volume groups, if applicable) 	Permits 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.
View diagnostic event log 	Starts the Event Log Viewer which displays a detailed list of events that occur in a disk subsystem.
Monitor performance 	Opens the Performance Monitor which provides information about how the disk subsystem is functioning.
Recover from failures 	Initiates the Recovery Guru which is used to help troubleshoot disk subsystem problems. Note If the disk subsystem is in a Needs Attention state  , the icon on the Recovery Guru toolbar button flashes.
Find node in tree 	Starts the Find dialog that allows you to search for a particular node in the Logical/Physical View or Mappings View of the Array Management Window.

Protecting Your Data

This chapter describes disk subsystem configuration options you can use to maximize data availability and to prevent data loss. A disk subsystem includes redundant drives, controllers, power supplies, and fans. This hardware redundancy keeps the disk subsystem working if a component fails. In addition, you can use the storage management software to implement the data protection options in this chapter.

For conceptual information and detailed procedures for the options described in this section, refer to *Learn About Data Protection Strategies in the Array Management Window* online help.

Software Redundancy

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

The following sections provide a brief overview of the premium features used for data protection strategies. For more detailed information, refer to [Chapter 4, “Configuring Disk Subsystems.”](#)

Remote Volume Mirroring

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

Snapshot Volumes

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

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

RAID Levels and Data Redundancy

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

RAID relies on a series of configurations, called levels, to determine how user and redundancy data is written and retrieved from the drives. Each level provides different performance and protection features. The storage management software offers four formal RAID level configurations: RAID levels 0, 1, 3, and 5. [Table 3-1](#) describes these configurations.

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 or an error-correcting code derived from the data. If a drive fails, the redundancy data can be used to quickly reconstruct information.

Only *one* RAID level can be configured across each volume group. Each volume group stores its own redundancy data. 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 data depends on the RAID level used.

Table 3-1 RAID Level Configurations (1 of 2)

RAID Level	Short Description	Detailed Description
RAID 0	Non-Redundant, Striping Mode	<ul style="list-style-type: none"> ● Used for high performance needs, but does not provide <i>data redundancy</i>. ● Stripes data across all drives in the volume group. ● Not recommended for high data availability needs. RAID 0 is better for non-critical data. ● A single drive failure causes all associated volumes to fail and data loss can occur.

Table 3-1 RAID Level Configurations (2 of 2)

RAID Level	Short Description	Detailed Description
RAID 1	Striping/Mirroring Mode	<ul style="list-style-type: none"> ● Also called RAID 10 or 0+1. ● Offers the best availability high performance; however, only half of the drives in the volume group are available for user data. ● Uses disk mirroring to make an exact copy from one drive to another drive. ● A single drive failure causes associated volumes become degraded, but the mirror drive allows access to the data. ● Can survive multiple drive failures as long as no more than one failure exists per mirrored pair. ● A drive-pair failure in a volume group causes all associated volumes to fail and data loss could occur.
RAID 3	High Bandwidth Mode	<ul style="list-style-type: none"> ● Both user data and redundancy data (parity) are striped across the drives. ● The equivalent of one drive's capacity is used for redundancy data. ● Good for large data transfers in applications such as multimedia or medical imaging that write and read large sequential chunks of data. ● A single drive failure in a volume group causes associated volumes to become degraded, but the redundancy data allows access to the data. ● Two or more drive failures in a volume group causes all associated volumes to fail and data loss could occur.
RAID 5	High I/O Mode	<ul style="list-style-type: none"> ● Both user data and redundancy data (parity) are striped across the drives. ● The equivalent of one drive's capacity is used for redundancy data. ● Good for multi-user environments such as database or file system storage, where typical input/output (I/O) size is small and there is a high proportion of read activity. ● A single drive failure in a volume group causes associated volumes to become degraded, but the redundancy data allows access to the data. ● Two or more drive failures in a volume group causes all associated volumes to fail and data loss could occur.

Hardware Redundancy

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

Controller Cache Memory

CAUTION **Sometimes write caching is disabled when batteries are low or discharged.** If a parameter called Write Caching Without Batteries is enabled on a volume, write caching continues even when batteries in the control module or array module are discharged. If you do not have an uninterruptible power supply (UPS) for power loss protection, do not enable this parameter, because data in the cache will be lost during a power outage if the control module or array module does not have working batteries.

Write caching can increase I/O performance during data transfers. However, it also increases the risk of data loss if a controller (or its memory) fails while unwritten data resides in cache memory. Write cache mirroring protects data during a controller or cache memory failure. When write cache mirroring is enabled, 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 alternate controller. Therefore, if one controller fails, the alternate can complete all outstanding write operations.

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

Low start and stop percentages provide for maximum data protection. However, in both cases, this increases the chance that data requested by a read command will not be in the cache, decreasing the cache hit percentage for writes and the I/O request. 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.

Data in the controller cache memory is also protected in case of power outages. Control modules and array modules contain batteries that protect the data in cache by keeping a level of power until the data can be written to the drive media. If a power outage occurs and there is no battery or the battery is damaged, 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. It is, therefore, important to change the control module and array module batteries at the recommended time intervals. The controllers in the disk subsystem keep track of the age (in days) of the battery. After replacing the battery, the age must be reset so that you will receive an accurate critical alert notification when the battery is nearing expiration and when it has expired.

Hot Spare Drives

A valuable strategy to protect data is to assign available drives in the disk subsystem as hot spare drives. A hot spare is a drive, containing no data, that acts as a standby in the disk subsystem in case a drive fails in a RAID 1, 3, or 5 volume. The hot spare adds another level of redundancy to the disk subsystem. Generally, the drive assigned as a hot spare should have a capacity that is equal to or greater than the capacity of the largest drive on the disk subsystem. If a drive fails in the disk subsystem, the hot spare is automatically substituted for the failed drive without requiring user intervention. If a hot spare is available when a drive fails, the controller uses redundancy data to reconstruct the data onto the hot spare. When you have physically replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This is called copyback.

If you do not have a hot spare, you can still replace a failed disk drive while the disk subsystem 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 onto the replacement drive. This is called reconstruction.

Background Media Scan

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

The media scan must be enabled for the entire disk subsystem as well as enabled on each volume.

The media scan runs on all volumes in the disk subsystem for which it has been enabled. The advantage of enabling a media scan 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, it also scans the volume redundancy data.

Channel Protection

In a SCSI environment, whether or not the volume group candidate has channel protection depends on the RAID level of the volume group and how many drives are present on any single drive channel. For example, a RAID 5 volume group would not have channel protection if more than one drive was present on a single drive channel.

In a Fibre Channel environment, channel protection is usually present for any volume group candidate because, when the disk subsystem is properly cabled, there are two redundant Fibre Channel Arbitrated Loops for each drive.

I/O Data Path Protection

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

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

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

For operating system-specific failover options, refer to the *BladeStore Storage Manager Installation Guide*.

Multi-Path Driver With AVT Enabled

If AVT is enabled when a volume is created, a controller must be assigned to own the volume (called the preferred controller, or preferred owner). The preferred controller normally receives the I/O requests to the volume. If a problem along the data path (such as a component failure) causes an I/O to fail, the multi-path driver will issue the I/O to the alternate controller.

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

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

Multi-Path Driver with AVT Disabled

When AVT is disabled, the I/O data path will still be protected as long as a multi-path driver is installed on each host connected to the disk subsystem. However, when an I/O request is sent to an individual volume, and a problem occurs along the data path to its preferred controller, *all* volumes on that controller will be transferred to an alternate controller instead of just that particular volume.

Password Protection

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

IMPORTANT If you forget the password, contact technical support.

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

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

Configuring Disk Subsystems

This chapter provides descriptions for volumes and volume groups, Dynamic Volume Expansion (DVE), and premium features such as SANshare Storage Partitioning, snapshot volumes, and Remote Volume Mirroring. Further outlined in this chapter are the specific functions of the Mappings View in the Array Management Window, and an overview of the heterogeneous host setting.

For additional conceptual information and detailed procedures for the options described in this section, refer to [Learn About Configuring a Disk Subsystem in the Array Management Window](#) online help.

Volumes and Volume Groups

When configuring a disk subsystem, appropriate data protection strategies as well as how the total storage capacity will be organized into volumes and shared among hosts must be considered.

The storage management software identifies several distinct volumes: standard, snapshot, snapshot repository, primary, secondary, and mirror repository.

- **Standard volume** – a logical structure created on a disk subsystem for data storage. A standard volume is created using the Create Volume Wizard. If the premium feature is not enabled for snapshot volumes or Remote Volume Mirroring, then only standard volumes will be created. Standard volumes are also used in conjunction with creating snapshot volumes and remote mirror volumes.
- **Snapshot volume** – a point-in-time image of a standard volume. A snapshot volume is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space. The volume from which you are basing the snapshot volume, called the base volume, must be a standard volume in your disk subsystem.
- **Snapshot repository volume** – a special volume in the disk subsystem created as a resource for a snapshot volume. A snapshot repository volume contains snapshot volume metadata and copy-on-write data for a particular snapshot volume.
- **Primary volume** – a standard volume in a mirror relationship that accepts host I/O and stores application data. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the associated secondary volume.
- **Secondary volume** – a standard volume in a mirror relationship that maintains a mirror (or copy) of the data from its associated primary volume. The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure of the primary site, the system administrator can promote the secondary volume to a primary role.
- **Mirror repository volume** – a special volume created as a resource for each controller in both the local and remote disk subsystem. The controller stores mirroring information on this volume, including information about remote writes that are not yet complete. The controller can use this information to recover from controller resets and accidental powering-down of disk subsystems.

Volume Groups

A volume group is a set of drives that the controller logically groups together to provide one or more volumes to an application host. When creating a volume from unconfigured capacity, the volume group and the volume are created at the same time. When creating a volume from free capacity, an additional volume is created on an existing volume group (Figure 4-1).

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

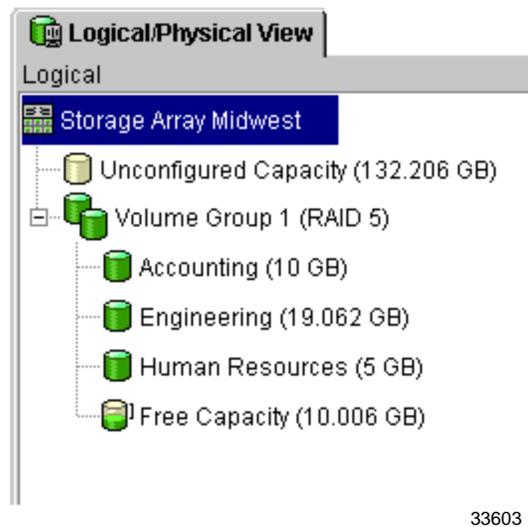


Figure 4-1 Unconfigured and Free Capacity Nodes

Standard Volumes

IMPORTANT The host operating system may have specific limits on how many volumes the host can access which must be considered when creating volumes for use by a particular host. For operating system restrictions, refer to the *BladeStore Storage Manager Product Release Notes* that were shipped with the software and your host operating system documentation.

A standard volume is a logical structure created on a disk subsystem for data storage. A volume is defined over a set of drives called a volume group, and has a defined RAID level and capacity. Volumes are created from either unconfigured capacity or free capacity nodes on the disk subsystem. If no volumes on the disk subsystem are configured, the only node available is the unconfigured capacity node ([Figure 4-1 on page 4-3](#)).

The Create Volume Wizard is used to create one or more volumes on the disk subsystem. During the volume creation process, the Wizard prompts you to select the capacity to allocate for the volumes and to define basic and optional advanced volume parameters for the volume. Each Wizard screen has context-sensitive help.

Specifying Volume Parameters from Free Capacity

IMPORTANT The free capacity, unconfigured capacity, or unassigned drives selected when starting the Wizard determine the default initial capacity selections. After the Wizard begins, the capacity can be changed by selecting a different free capacity node location for the volume, or by selecting different unassigned drives for the volume group.

The Specify Volume Parameters screen is used to specify the capacity for the volume, the volume name, and whether to use recommended advanced volume settings or customize the advanced volume properties for this volume.

The Specify Advanced Volume Parameters screen is also used to specify the volume I/O characteristics that will apply to the volume based on the needs of your application, or a custom cache read-ahead multiplier and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter. Volumes are tailored to specific application needs by customizing the advanced volume settings.

After the volume creation process is finished, a confirmation dialog is displayed. Use this dialog to restart the Wizard to create another volume.

Specifying Volume Parameters from Unconfigured Capacity

IMPORTANT The free capacity, unconfigured capacity, or unassigned drives selected when starting the Wizard determine the default initial capacity selections. After the Wizard begins, the capacity can be changed by selecting a different free capacity node location for the volume, or by selecting different unassigned drives for the volume group.

The Specify Volume Group Parameters screen is used to specify the RAID level of the volume group to meet your volume data storage and protection requirements, and to select the drives that will comprise the volume group. It also provides a way to specify the capacity for the volume, the volume name, and whether to use recommended advanced volume settings or customize the advanced volume properties for the volume.

The Specify Advanced Volume Parameters screen can further be used to specify the volume I/O characteristics that will apply to the volume based on the needs of your application, including a custom cache read-ahead multiplier and segment size, preferred controller ownership, and a volume-to-LUN mapping parameter.

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

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

After the volume creation process is finished, a confirmation dialog is displayed. Use this dialog to restart the Wizard to create another volume.

Snapshot Volumes

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

A snapshot volume is typically created so that an application, such as a backup application, can access the snapshot volume and read the data while the base volume remains online and user-accessible. When the backup completes, the snapshot volume is no longer needed. You can also create several snapshot volumes of a base volume and write data to the snapshot volumes to perform testing and analysis. Before upgrading a database management system, for example, this feature can be used to test different configurations. Then the performance data provided by the storage management software can be used to help decide how to configure the live database system. The maximum number of snapshot volumes allowed is one half of the total volumes supported by your controller model, while the maximum number of snapshot volumes supported by a volume is four.

When a snapshot volume is created, the controller suspends I/O to the base volume for a few seconds while it creates a physical volume, called the snapshot repository volume, to store the snapshot volume metadata and copy-on-write data. Because the only data blocks that are physically stored in the snapshot repository volume are those that have changed since the time the snapshot volume was created, the snapshot volume uses less disk space than a full physical copy.

The storage management software provides a warning message when the snapshot repository volume nears a user-specified threshold (a percentage of its full capacity, the default is 50%). When this condition occurs, the capacity of the snapshot repository volume can be expanded from free capacity on the volume group. If you are out of free capacity on the volume group, unconfigured capacity can be added to the volume group to expand the snapshot repository volume.

The Create Snapshot Volume Wizard is used to create snapshot volumes by defining the snapshot volume's name, the name of its associated snapshot repository volume, and to specify the snapshot repository volume's capacity as a percentage of the base volume's capacity. Either the Simple or Advanced path can be selected through the Create Snapshot Volume Wizard:

- **Simple Path** – provides a preview of the preconfigured snapshot volume and snapshot repository volume parameters.
- **Advanced Path** – provides a way to select a free capacity or unconfigured capacity node on which to place the snapshot repository volume, and allows you to change snapshot repository volume parameters. The Advanced Path can be chosen regardless of whether free capacity or unconfigured capacity is selected.

Disabling or Deleting a Snapshot Volume

As long as a snapshot volume is enabled, disk subsystem performance is impacted by the copy-on-write activity to the associated snapshot repository volume. If a snapshot volume is no longer needed, it can be disabled, which will stop the copy-on-write activity.

If the snapshot volume is disabled, you can retain it and its associated snapshot repository volume. When you need to create a different point-in-time image of the same base volume, you can use the re-create option to reuse the disabled snapshot volume and its associated snapshot repository volume. This takes less time than creating a new one.

When you disable a snapshot volume:

- You cannot use the snapshot volume again until you use the re-create option.
- Only the snapshot volume is disabled. All other snapshot volumes remain functional.

If you do not intend to re-create a snapshot volume, you can delete the snapshot volume instead of disabling it. When you delete a snapshot volume, the associated snapshot repository volume is also deleted.

Re-creating a Snapshot Volume

If you have a snapshot volume that you no longer need, instead of deleting it, you can reuse it (and its associated repository volume) to create a different point-in-time image of the same base volume. Re-creating a snapshot volume takes less time than creating a new one.

When you re-create a snapshot volume:

- The snapshot volume must have either an optimal or a disabled state.
- All copy-on-write data previously on the snapshot repository volume is deleted.
- Snapshot volume and snapshot repository volume parameters remain the same as the previously disabled snapshot volume and its associated snapshot repository volume. After the snapshot volume is re-created, you can change parameters on the snapshot repository volume through the appropriate menu options.
- The original names for the snapshot volume and snapshot repository volumes will be retained. You can change these names after the re-create option is completed.

Remote Volume Mirroring

This is a premium feature of the storage management software and must be enabled either by you or your storage vendor. The Remote Volume Mirroring feature is used for online, real-time replication of data between disk subsystems over a remote distance.

When you create a remote volume mirror, a mirrored volume pair is created and consists of a primary volume at a primary disk subsystem and a secondary volume at a secondary disk subsystem. Prior to creating a mirror relationship, the Remote Volume Mirroring feature must be enabled and activated on both the primary and secondary disk subsystems. The primary volume is the volume that accepts host I/O and stores application data. When the mirror relationship is first created, data from the primary volume is copied in its entirety to the secondary volume. The secondary volume maintains a mirror (or copy) of the data from its associated primary volume. The secondary volume remains unavailable to host applications while mirroring is underway. In the event of a disaster or catastrophic failure of the primary site, the secondary volume can be promoted to a primary role.

Mirror Relationships

A secondary volume candidate must be created on the secondary site if one does not already exist and must be a standard volume of equal or greater capacity than the associated primary volume.

When a secondary volume candidate is available, a mirror relationship can be established in the storage management software by identifying the disk subsystem containing the primary volume and the disk subsystem containing the secondary volume.

Mirror Repository Volumes

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

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

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

Data Replication

Data replication between the primary volume and the secondary volume is managed by the controllers and is transparent to host machines and applications.

When the controller owner of the primary volume receives a write request from a host, the controller first logs information about the write to a mirror repository volume, then writes the data to the primary volume. The controller then initiates a remote write operation to copy the affected data blocks to the secondary volume at the secondary disk subsystem.

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

Register the Volume with the Operating System

IMPORTANT The `hot_add` utility is not available for all operating systems. Consult the *BladeStore Storage Manager Installation Guide* for information on operating system availability and how to run this utility.

After creating all volumes and assigning volume-to-LUN mappings, the host-based `hot_add` utility is used to register the volume with the operating system. The host-based `SMdevices` utility (if applicable for your operating system) is used to associate the mapping between the physical device name and the volume name. Refer to [“SANshare Storage Partitioning” on page 4-14](#) for more information.

Dynamic Volume Expansion (DVE)

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

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

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

You cannot add storage capacity to a volume if:

- One or more hot spare drives are in use in the volume
- The volume has a non-optimal status
- Any volume in the volume group is in any state of modification
- The controller that owns this volume is in the process of adding capacity to another volume (Each controller can add capacity to only one volume at a time.)
- No free capacity exists in the volume group, *and* no unconfigured capacity is available to add to the volume group
- No unconfigured capacity (in the form of drives) is available to add to the volume group

Mappings View

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

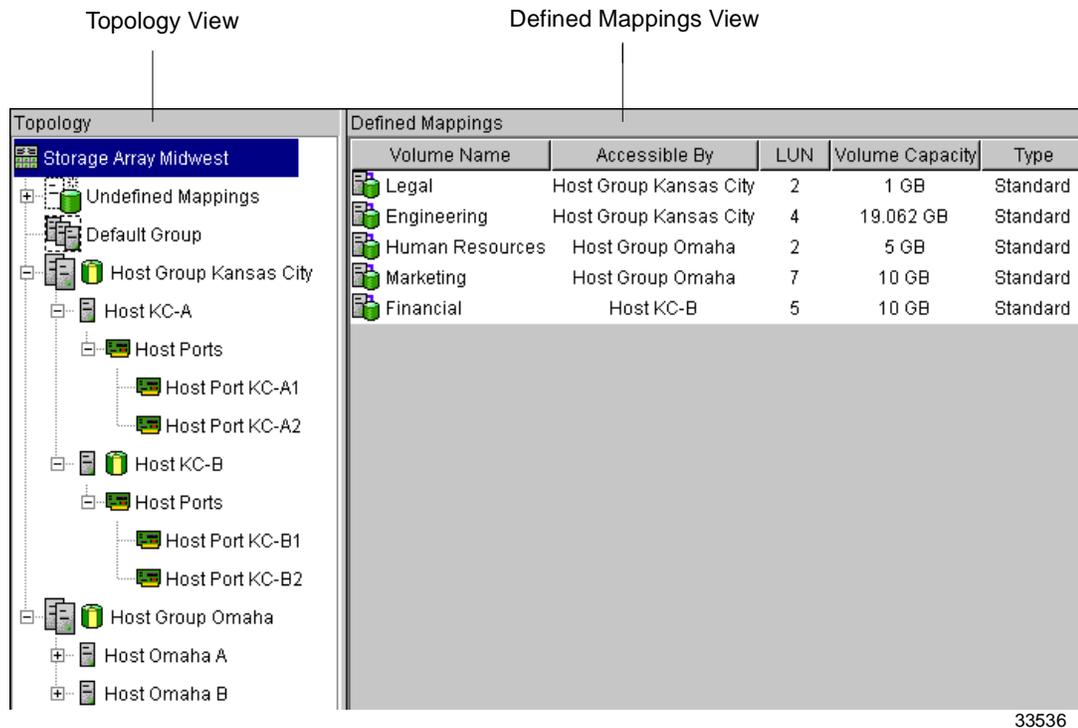


Figure 4-2 Mappings View Window

Table 4-1 Mappings View Tab

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

Table 4-2 describes the topological elements displayed in the Mappings View Window in Figure 4-2 on page 4-11.

Table 4-2 Volume-to-LUN Terminology (1 of 2)

Term	Description
SANshare Storage Partitioning Topology	<p>A collection of nodes (default group, host groups, hosts, and host ports) shown in the Topology View of the Mappings View tab.</p> <p>You must define the various topological elements if you want to define specific volume-to-LUN mappings and storage partitions for host groups or hosts.</p>
Default Group	<p>A node in the Topology View that designates all host groups, hosts, and host ports that: (1) have no specific volume-to-LUN mappings and (2) share access to any volumes that were automatically assigned default LUN mappings by the controller firmware during volume creation.</p>
Host Group	<p>An optional topological element that you define if you want to designate a collection of hosts that will share access to the same volumes. The host group is a logical entity.</p>
Host	<p>A computer that is attached to the disk subsystem and accesses various volumes on the disk subsystem through its host ports (host bus adapters). You can define specific volume-to-LUN mappings to an individual host or assign the host to a host group that shares access to one or more volumes.</p>
Host Port	<p>The physical connection that allows a host to gain access to the volumes in the disk subsystem. When the host bus adapter only has one physical connection (host port), the terms host port and host bus adapter are synonymous.</p> <p>Host ports can be automatically detected by the storage management software after the disk subsystem has been connected and powered-up. Therefore, if you want to define specific volume-to-LUN mappings for a particular host or create storage partitions, you must define the host's associated host ports.</p> <p>Initially, all detected host ports belong to the Default Group. Therefore, if during volume creation, you had a LUN automatically assigned to a volume, that volume will be accessible by any of the host ports in the Default Group. If you have the SANshare Storage Partitioning feature enabled, then you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation.</p> <p>Use the host bus adapter utility to find the World Wide Name (WWN) of the host port. (This is the host port identifier shown in the Define New Host Port dialog in the Mappings View.) The WWNs for the host ports on a particular host are used to define the host ports and associate them with a particular host using the Define New Host Port dialog. If necessary, refer to your operating system or host bus adapter documentation for more information.</p> <p>If a host port is moved, any volume-to-LUN mappings must be re-mapped. Access to your data will be lost until this is done.</p>

Table 4-2 Volume-to-LUN Terminology (2 of 2)

Term	Description
Logical Unit Number (LUN)	<p>The number a host uses to access a volume on a disk subsystem. Each host has its own LUN address space. Therefore, the same LUN may be used by different hosts to access different volumes on the disk subsystem. However, a volume can only be mapped to a single LUN. A volume cannot be mapped to more than one host group or host.</p> <p>For example, Figure 4-5 on page 4-18 shows that Host Group Kansas City may access Volume Legal using LUN 2 and Host Group Omaha may access Volume HResources also using LUN 2.</p>
Default Volume-to-LUN mapping	<p>During volume creation, you can specify that you want to have the software assign a LUN automatically to the volume or that you want to map a LUN to the volume later. If you have the SANshare Storage Partitioning feature enabled, than you should always choose to map the volume later using the options in the Mappings View so that a LUN is not automatically assigned to a volume during volume creation. Any volumes that are given automatic (default) volume-to-LUN mappings can be accessed by all host groups or hosts that do not have specific volume-to-LUN mappings. These host groups and hosts are shown as part of the Default Group in the Topology section of the Mappings View.</p>
Specific Volume-to-LUN mapping	<p>A specific volume-to-LUN mapping occurs when you select a defined host group or host in the Topology View and select the SANshare Storage Partitioning Wizard or Define Additional Mapping option to assign a volume a specific LUN (volume-to-LUN mapping). This designates that only the selected host group or host has access to that particular volume through the assigned LUN. You can define one or more specific volume-to-LUN mappings for a host group or host.</p> <p>Note The SANshare Storage Partitioning feature must be enabled to create specific mappings.</p>

SANshare Storage Partitioning

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

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

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

A maximum of 64 storage partitions can be supported by the storage management software. The software can further support up to two host ports in each host and up to eight ports in each host group, allowing a four-way cluster of dual-adapter hosts. After the total storage capacity is configured into volumes, those volumes need to be shared among hosts in the disk subsystem. SANshare Storage Partitioning allows hosts with the same or different operating systems (heterogeneous hosts) to share access to a disk subsystem.

SANshare Storage Partitioning Example

In the example shown in [Figure 4-3](#), four hosts (Omaha A and B and KC-A and B) are connected to Disk Subsystem Midwest. Host KC-B has exclusive access to Volume Financial. Hosts KC-A and KC-B (Host Group Kansas City) share access to Volumes Legal and Engineering, and Omaha A and Omaha B (Host Group Omaha) share access to Volumes Marketing and HResources. So, three storage partitions were configured.

The first partition is composed of Volume Financial. This volume is accessed by Host KC-B using LUN 5. Even though Host KC-B is part of the logical Host Group Kansas City, Host KC-A cannot access this volume because the volume-to-LUN mapping was created with Host KC-B rather than the Host Group Kansas City.

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

The third partition consists of Volumes Marketing and 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 using both LUNs 7 and 2.

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

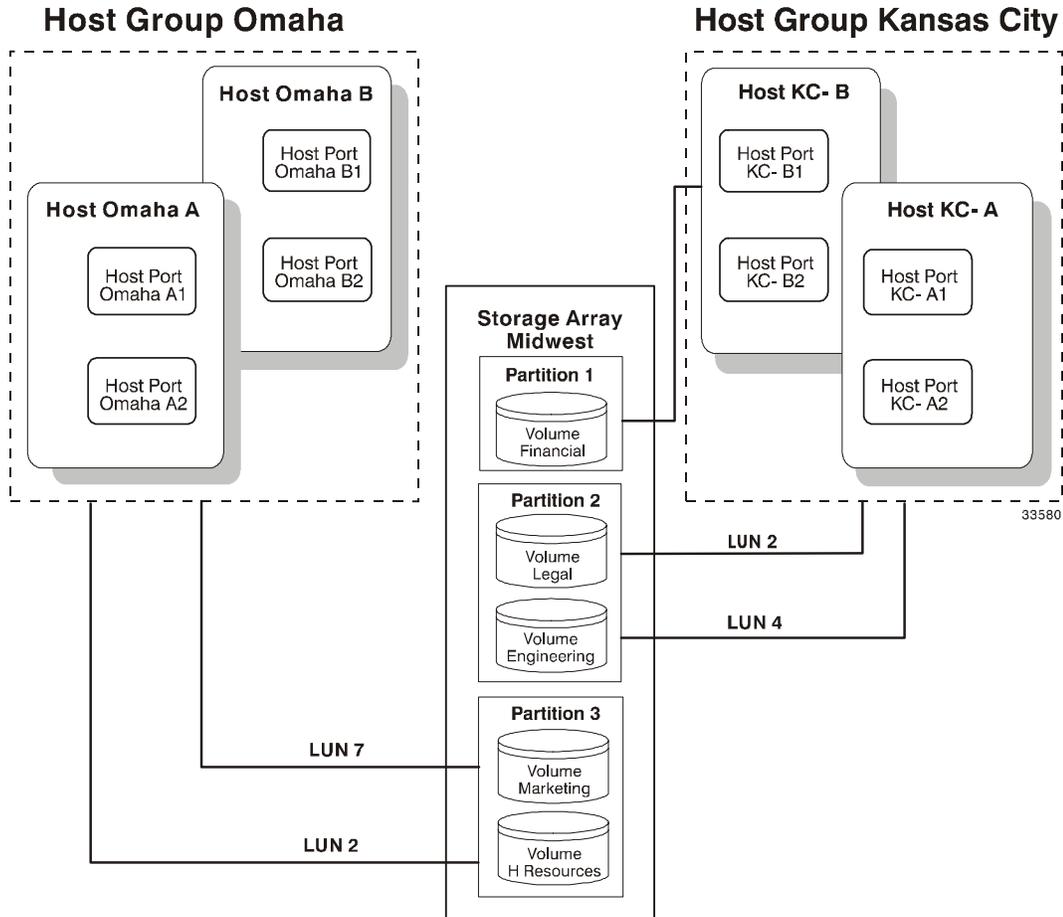


Figure 4-3 SANshare Storage Partitioning Example

SANshare Storage Partitioning involves three key steps:

- Create volumes on the disk subsystem. As part of the volume creation, specify one of two volume-to-LUN mapping settings:
 - **Automatic** – If you are *not* using SANshare Storage Partitioning, specify this setting. The Automatic setting specifies that a LUN be automatically assigned to the volume using the next available LUN within the Default Group. This setting will grant volume access to host groups or hosts that have no specific volume-to-LUN mappings (designated by the Default Group in the Topology View).
 - **Map later with SANshare Storage Partitioning** – If you *are* using SANshare Storage Partitioning, specify this setting. The Map later setting specifies that a LUN not be assigned to the volume during volume creation. This setting allows definition of a specific volume-to-LUN mapping and creation of storage partitions.
- Define the storage partition topology (including host groups, hosts, and host ports) that access the volumes. Storage partition topology is reconfigurable. You can:
 - Move a host port
 - Replace a host port
 - Move a host from one host group into another host group
 - Delete a host group, host, or host port
 - Rename a host group, host, or host port
 - Change a volume-to-LUN mapping
 - Define additional volume-to-LUN mappings
- Grant volume access to defined host groups or hosts by defining volume-to-LUN mappings, using the SANshare Storage Partitioning Wizard. Each host group or host is granted a unique view of partitioned storage.

A defined host group or host can either access:

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

Heterogeneous Hosts

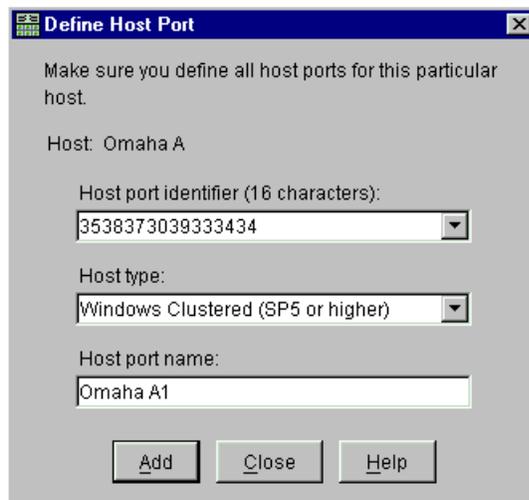
The heterogeneous hosts portion of the SANshare Storage Partitioning feature allows hosts running different operating systems to access a single disk subsystem. To specify different operating systems for attached hosts, you must specify the appropriate host type when you define the host ports for each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

Host types can be completely different operating systems, such as Solaris and Windows NT, or variants of the same operating system, such as Windows NT – clustered and Windows NT – non-clustered. When a host type is specified, it allows the controllers in the disk subsystem to tailor their behavior (such as LUN reporting and error conditions) depending on the operating system or variant of the host sending the information.

Heterogeneous Hosts Example

IMPORTANT Heterogeneous host settings are only available with SANshare Storage Partitioning enabled. In [Figure 4-5 on page 4-18](#), the SANshare Storage Partitioning feature is enabled.

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



63076

Figure 4-4 Host Port Definitions Dialog

In the example shown in [Figure 4-5](#), heterogeneous hosts have been defined on Disk Subsystem Midwest. Four hosts (Omaha A and B and KC-A and B) are connected to Disk Subsystem Midwest. Hosts Omaha A and Omaha B share access to Volumes Marketing and HResources. Host KC-A has exclusive access to Volumes Legal and Engineering, and Host KC-B has exclusive access to Volume Financial. So, three storage partitions have been configured. Because there are four hosts running three different operating systems, the appropriate host types must be defined for each host port to support these heterogeneous hosts.

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

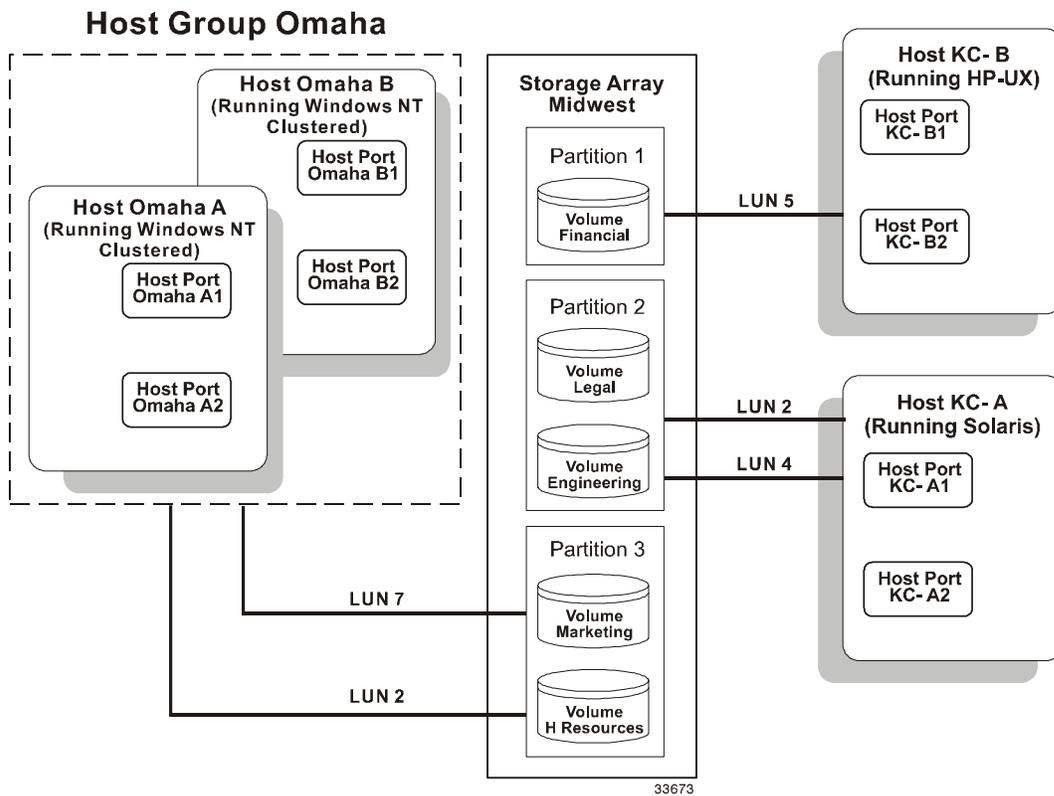


Figure 4-5 Heterogeneous Hosts Example

Maintaining and Monitoring Disk Subsystems

This chapter describes methods for maintaining disk subsystems in the management domain, including troubleshooting disk subsystem problems, recovering from a disk subsystem problem using the Recovery Guru, and configuring alert notifications using the event monitor.

For additional conceptual information and detailed procedures for the options described in this section, refer to *Learn About Monitoring Disk Subsystems in the Enterprise Management Window* online help.

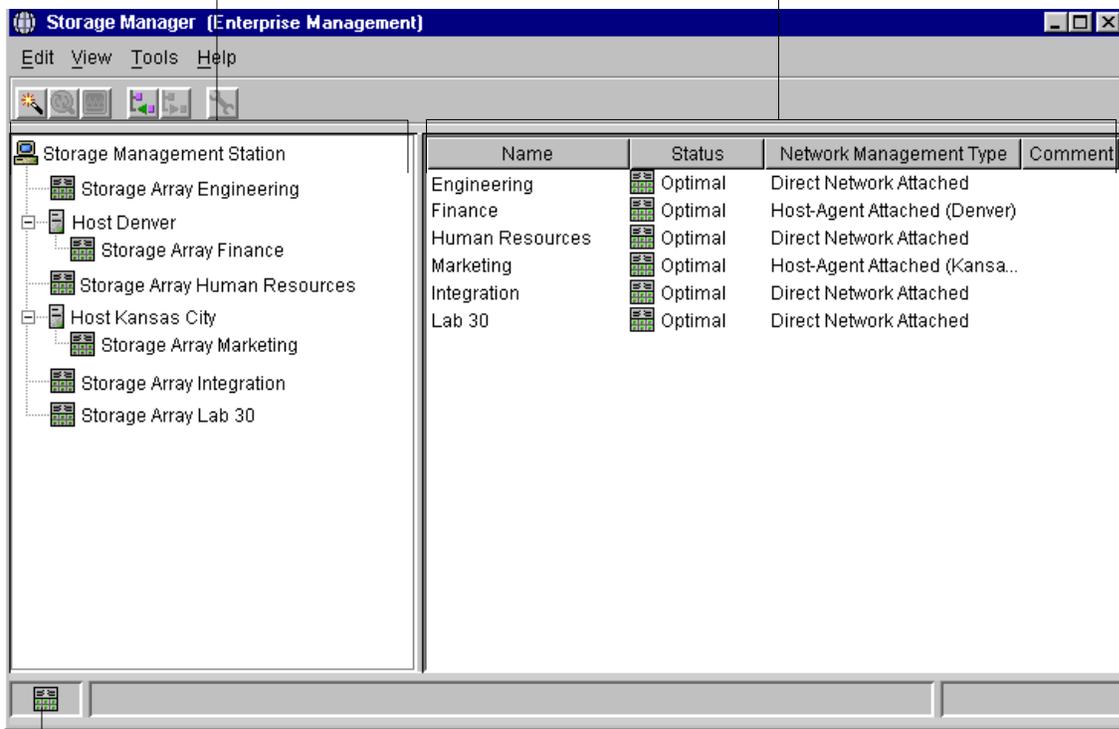
Disk Subsystem Health

IMPORTANT The Enterprise Management Window or the event monitor must be running to receive notification of critical events for the disk subsystems. In addition, alert notifications must be configured in the Enterprise Management Window.

The Enterprise Management Window provides a summary of the conditions of all known disk subsystems in the management domain. Appropriate status indicators will be shown in the Device Tree, the Device Table, and in the health summary status area in the lower-left corner of the Enterprise Management Window (Figure 5-1).

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

The *Device Table* contains a status column that displays the status of each disk subsystem.



63075

The *Overall Health Status* is a consolidated health status for all of the known disk subsystems in the Management Domain.

Figure 5-1 Monitoring Disk Subsystem Health Using the Enterprise Management Window

Disk Subsystem Status Icons

Table 5-1 provides information about the disk subsystem status icons that display:

- In the Device Tree, Device Table, and Overall Health Status area in the Enterprise Management Window
- As the Root Node in the Logical View Tree in the Array Management Window

Table 5-1 Disk Subsystem Status Icon Quick Reference

Status	Description
Optimal 	Indicates every component in the disk subsystem is in the desired working condition.
Needs Attention 	Specifies a problem on a disk subsystem that requires intervention to correct. To correct the problem, start the Array Management Window for that particular disk subsystem, and then use Recovery Guru to pinpoint the cause of the problem and obtain appropriate procedures.
Fixing 	Signifies a Needs Attention condition has been corrected and the disk subsystem is transitioning to an Optimal status; 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 disk subsystem status to change directly from Needs Attention to Optimal, without an interim status of Fixing. This icon is not displayed in the Overall Health Status area. The Optimal status icon is displayed instead.
Unresponsive 	Means the storage management station cannot communicate with the only controller or both controllers in the disk subsystem over its network management connection. Note This icon is not displayed in the Logical View of the Array Management Window. If the Array Management Window is open and the disk subsystem becomes Unresponsive, the last known status icon (Optimal, Needs Attention, or Fixing) is shown.
Contacting Device 	Designates that you have started the Enterprise Management Window and the storage management software is establishing contact with the disk subsystem. Note This icon is not displayed in the Logical View of the Array Management Window.

Event Monitor

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

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

Figure 5-2 shows how the event monitor and the Enterprise Management Window client software send alerts to a remote system. The storage management station contains a file with the name of the disk subsystem being monitored and the address where alerts will be sent. The alerts and errors that occur on the disk subsystem are continuously being monitored by the client software and the event monitor. The event monitor takes over for the client after the client software package is shut down. When an event is detected, a notification is sent to the remote system.

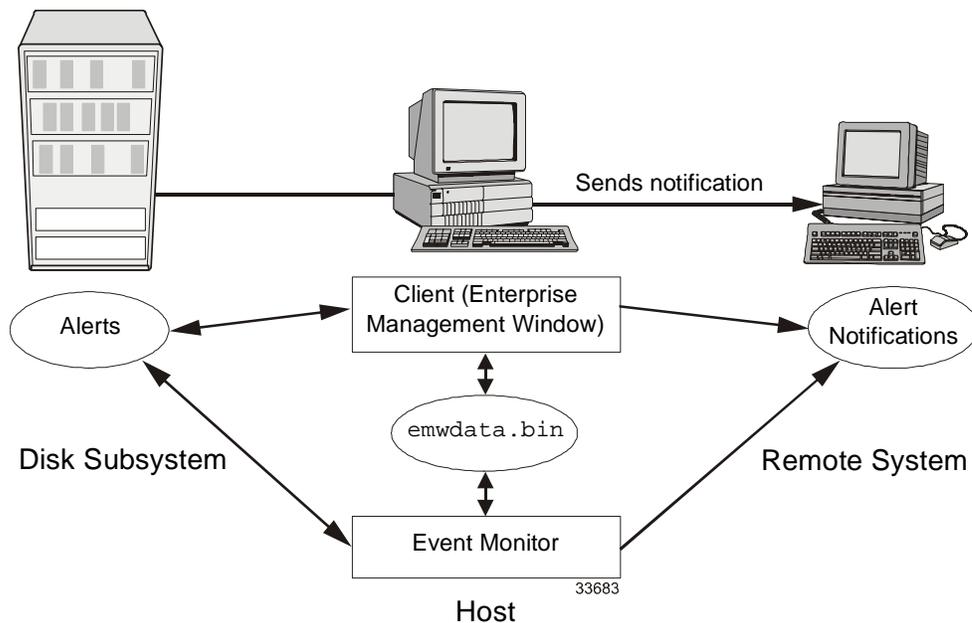


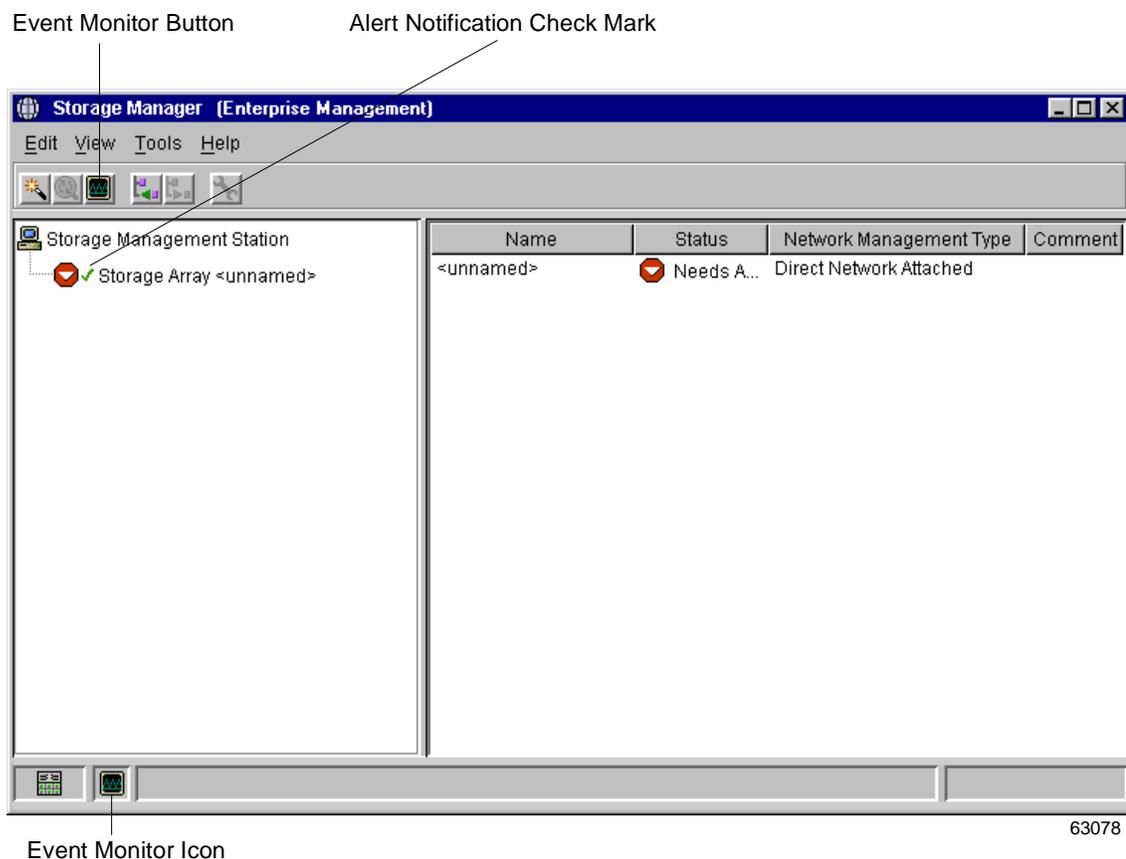
Figure 5-2 Event Monitor Configuration

Because the event monitor and the Enterprise Management Window share the information to send alert messages, the Enterprise Management Window contains some visual cues to assist in the event monitor installation and synchronization. The parts of the Enterprise Management Window that are related to event monitoring are shown in Figure 5-3.

Using the event monitor involves three key steps:

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

You must have administrative permissions to install software on the computer where the event monitor will reside. After the storage management software has been installed, the icon shown in Figure 5-3 will be present in the lower-left corner of the Enterprise Management Window.



63078

Figure 5-3 Event Monitor Example

- Set up the alert destinations for the disk subsystems you want to monitor from the Enterprise Management Window. A check mark indicates where the alert is set (storage management station, host, or disk subsystem). When a critical problem occurs on the disk subsystem, the event monitor will send a notification to the appropriate alert destinations that were specified.
- Synchronize the Enterprise Management Window and the event monitor. After it has been installed, the event monitor continues to monitor disk subsystems and send alerts as long as it continues to run. If you make a configuration change in the Enterprise Management Window, such as adding or removing a disk subsystem or setting additional alert destinations, you should manually synchronize the Enterprise Management Window and the event monitor.

Alert Notifications

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

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

IMPORTANT To set up alert notifications using SNMP traps, you must copy and compile a management information base (MIB) file on the designated network management station. Refer to the *BladeStore Storage Manager Installation Guide* for more information.

There are three key steps involved in configuring alert notifications:

- Select a node in the Enterprise Management Window that will display alert notifications for the disk subsystems you want to monitor. You can set the alert notifications at any level:
 - Every disk subsystem in the management domain
 - Every disk subsystem attached and managed through a particular host
 - An individual disk subsystem

- Configure e-mail destinations, if desired. You must provide a mail server name and an e-mail sender address for the e-mail addresses to work.
- Configure SNMP trap destinations, if desired. The SNMP trap destination is the IP address or the host name of a station running an SNMP service, such as a Network Management Station.

Customer Support Alert Notifications

IMPORTANT If you do *not* configure the customer support alert notifications option, the e-mail alert notification will contain only a summary of the critical event. If you *do* configure this option, all specified e-mail addresses will receive the summary, detailed information about the affected disk subsystem, and the specified contact information.

The Enterprise Management Window contains options to configure the system to send e-mail notifications to a specified customer support group if a critical event occurs on a disk subsystem. After it is configured, the e-mail alert notification includes a summary of the critical event, details about the affected disk subsystem, and customer contact information. Contact technical support for more information about setting up this file.

Configuring customer support alert notifications involves the following:

- 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 system administrators.
- Name the file `userdata.txt` and save it in the home directory (for example, `Winnt\profiles\`) on the client machine you are using to manage the disk subsystem (This may be your host machine if you installed the client software on the host.).
- Configure the alert notifications using e-mail or SNMP trap destinations.

Problem Notification

IMPORTANT The Enterprise Management Window or the event monitor must be running to receive notification of critical events for the disk subsystems. In addition, you must have configured the alert notifications in the Enterprise Management Window.

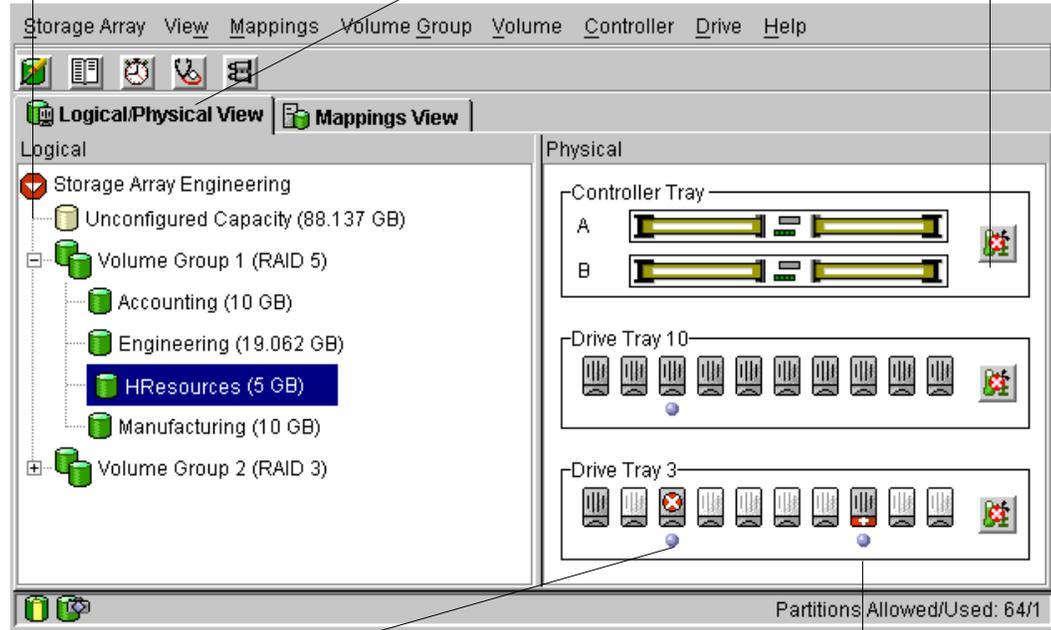
Use Recovery Guru to help troubleshoot disk subsystem problems. Where necessary, use the hardware documentation in conjunction with the recovery steps to replace failed components.

Typically, disk subsystem problems are indicated by:

- A Needs Attention status icon displayed in:
 - The Overall Health Status area, Device Tree View, and Device Table of the Enterprise Management Window.
 - The Array Management Window Logical View.
- The Recovery Guru Optimal 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.
- Receipt of critical SNMP or e-mail notifications.
- The hardware fault lights display.

In [Figure 5-4](#), the Array Management Window for Disk Subsystem Engineering indicates a components problem and a Needs Attention status in the Logical/Physical View.

The disk subsystem icon indicates a Needs Attention status. The Recovery Guru toolbar button is flashing. The Components button in the control module is indicating a failure.



63079

The Drive icon is indicating a failure.

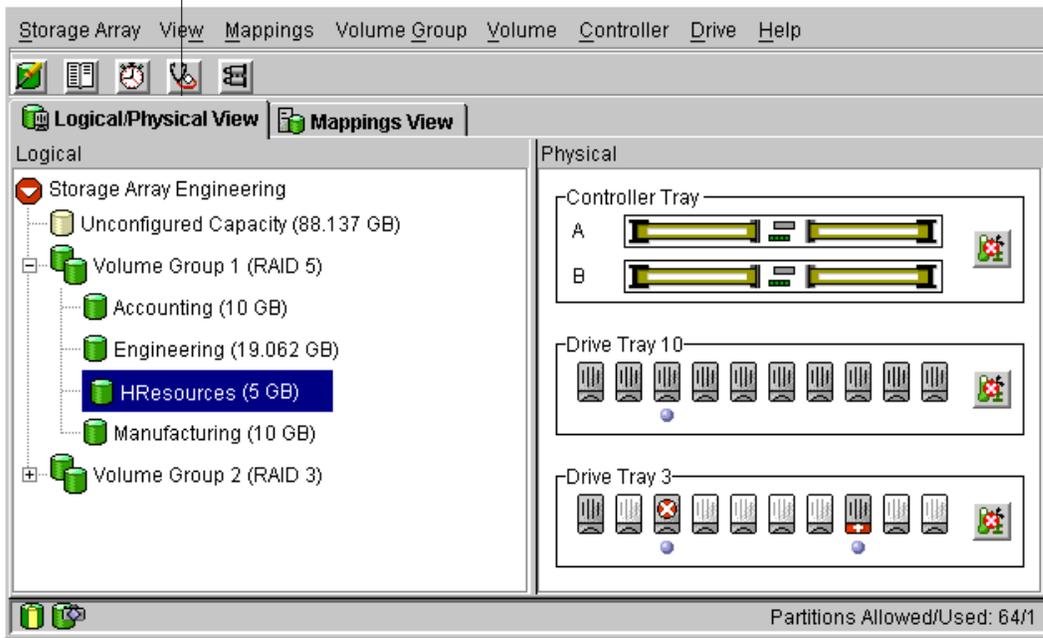
The hot spare drive is In Use (taking over for a failed drive). This means that the data from the failed drive has been reconstructed to the hot spare drive. Because the hot spare drive has taken over for the failed drive, the associated volume (Volume HResources) in Volume Group 1 remains Optimal.

Figure 5-4 Problem Notification in the Array Management Window

Disk Subsystem Problem Recovery

When you suspect a disk subsystem problem, launch the Recovery Guru. The Recovery Guru is a component of the Array Management Window that will diagnose the problem and provide the appropriate procedure to use for recovery. The Recovery Guru can be displayed by selecting the Recovery Guru toolbar button in the Array Management Window, shown in [Figure 5-5](#).

The Recovery Guru Toolbar Button



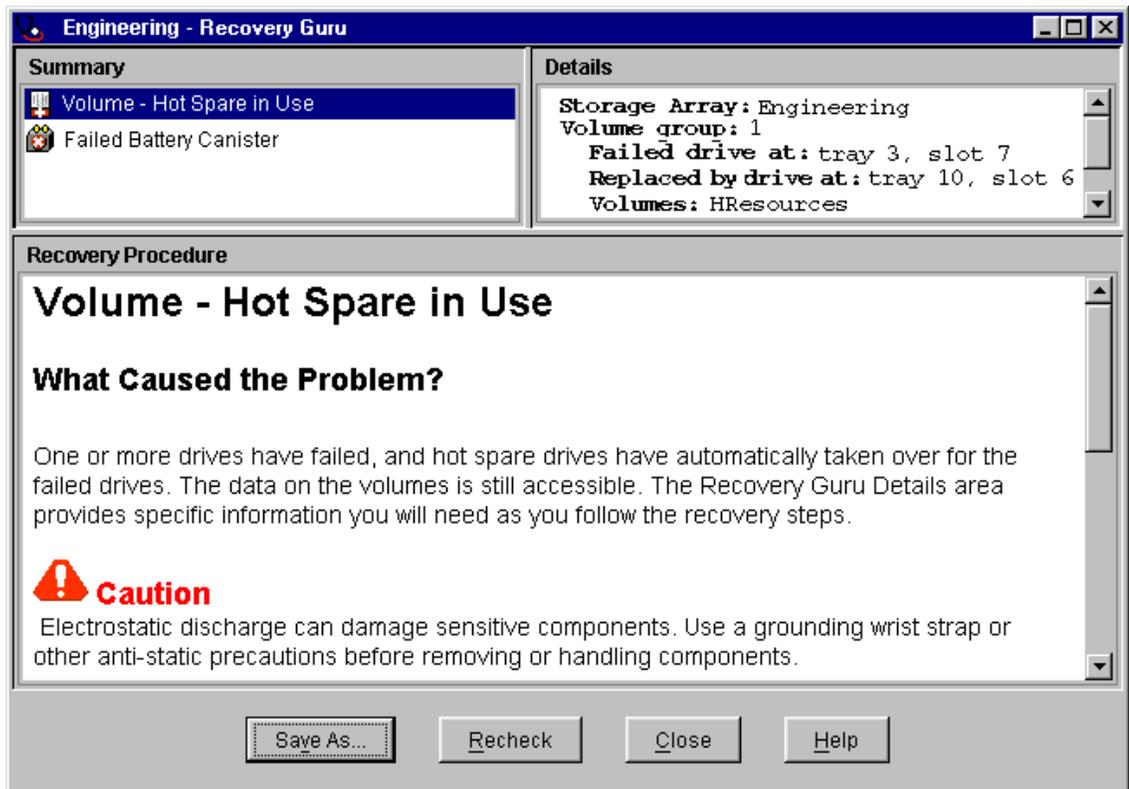
63079

Figure 5-5 Displaying the Recovery Guru Window

Recovery Guru Example

The Recovery Guru window is divided into three views: Summary, Details, and Recovery Procedures. The Summary view presents a list of disk subsystem problems. The Details view displays information about the selected problem in the Summary area. The Recovery Procedure view lists the appropriate steps to follow for the selected problem in the Summary view.

For example, in [Figure 5-6](#), the Summary area displays two different failures in this disk subsystem; a hot spare in use and a failed battery CRU. The Details area shows that in Volume HResources, a hot spare drive in tray 10, slot 6 has replaced a failed drive in tray 3, slot 7. The Recovery Procedure window explains the cause of the selected problem (seen in the Summary view), and describes the appropriate procedures needed to recover from this failure.

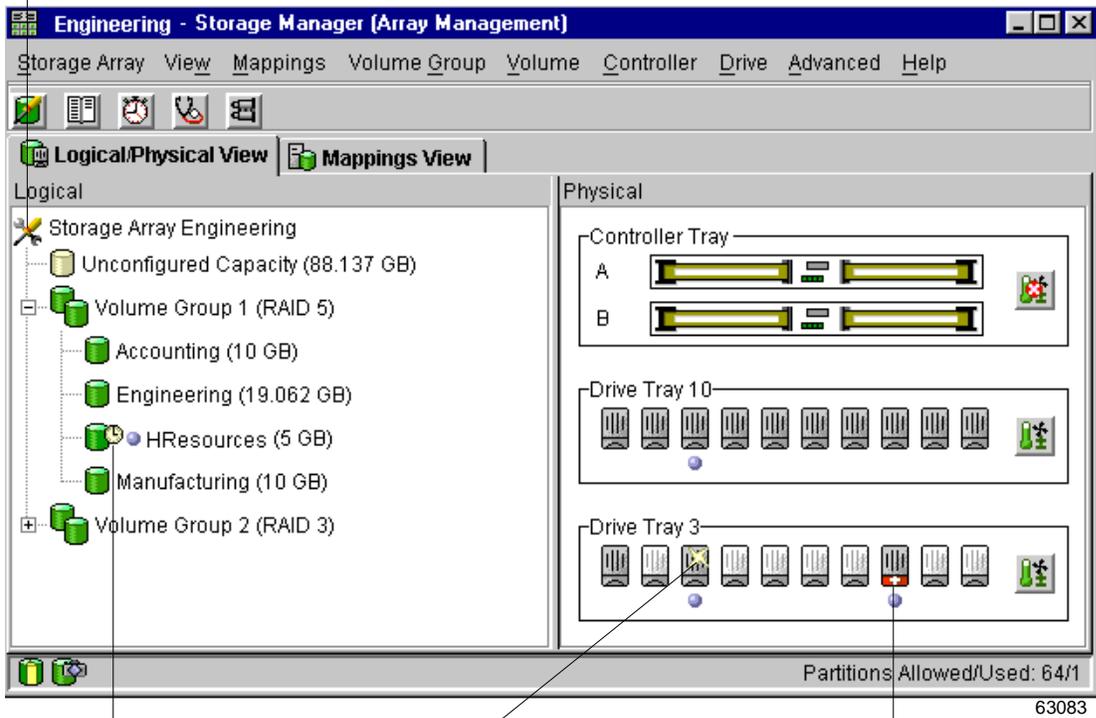


63082

Figure 5-6 Recovery Guru Window Example

As you follow the recovery procedure to replace the failed drive, the disk subsystem status icon changes to Fixing, the associated volume (HResources) status icon changes to Degraded–Modification in Progress, and the replaced drive status icon changes to Replaced. The data that was reconstructed to the hot spare drive is now being copied back to the replaced drive. These changes are shown in Figure 5-7.

The disk subsystem status icon changes from Needs Attention to Fixing.



Volume status icon changes from Optimal to Degraded-Modification in Progress.

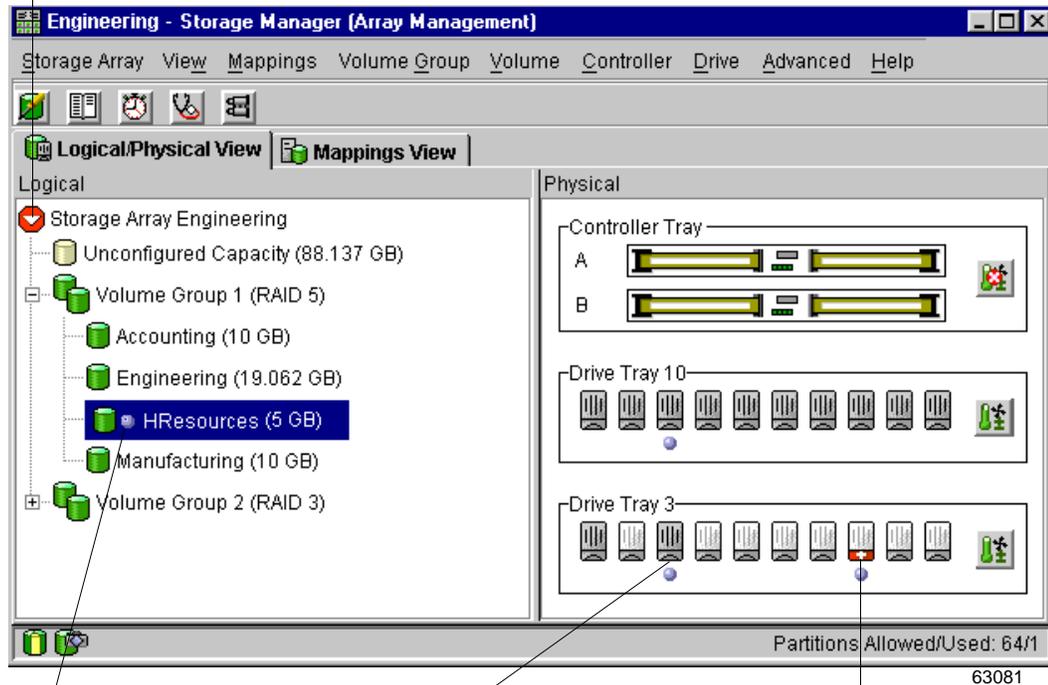
Drive status icon changes from Failed to Replaced.

Hot spare drive status icon remains Optimal, In Use during the copyback operation.

Figure 5-7 Status Icon Changes During an 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 5-8](#).

Disk subsystem status icon changes from Fixing back to Needs Attention.



Volume status icon changes from Operation in Progress back to Optimal.

Drive status icon changes from Replaced to Optimal, Assigned.

Drive status icon for the hot spare changes from Hot Spare - In Use to Hot Spare - Standby.

Figure 5-8 Status Changes When The Example Recovery Operation is Completed

After you replace the failed drive in the disk array:

- The disk subsystem status icon in the Logical View returns to Optimal.
- The disk subsystem status icon in the Enterprise Management Window changes from Needs Attention to Optimal.
- The Recovery Guru button stops blinking.

NOTE For the Recovery Guru button to register Optimal status, the failed battery must be replaced as well.

Index

A

- Alert Notifications, settings 5-6
- Array Management Window
 - figure 2-10
 - functions 2-9
 - introduction to 2-2
 - menu options 2-12
 - problem notification in 5-8
 - tabs 2-10
 - toolbar 2-13
- AVT (Auto-Volume Transfer) and RDAC, related to I/O 3-8

B

- batteries
 - control module 3-6
 - write caching 3-5

C

- channel protection 3-7
- client software, windows 2-2
- Components button 2-11
- Contacting Device status 5-3
- controller 4-12
- controller cache memory 3-5
- copyback 3-6
- Customer Support Alert Notification, settings 5-7

D

- data protection strategies
 - Hardware Redundancy 3-5
 - I/O (Input/Output) 3-8
 - RAID (Redundant Array of Independent Disks) 3-3
 - Software Redundancy 3-2
- Default Group, defined as a topological element 4-12
- default LUN 4-12

- Defined Mappings View 4-11
- Device Table 2-6, 5-3
- Device Tree 2-6, 5-3
- disable snapshot volumes 4-7
- disk subsystem health, summary of conditions 5-2
- DVE (Dynamic Volume Expansion), also called increase volume capacity 4-10

E

- Enterprise Management Window
 - figure 2-5
 - functions 2-4
 - introduction to 2-2
 - menu options 2-7
 - status icons displayed in 5-3
 - toolbar 2-8
- Event Monitor
 - overview 5-4
 - setting up alert destinations 5-6
 - synchronizing with Enterprise Management Window 5-6

F

- firmware versions supported 2-9
- Fixing status 5-3
- free capacity 4-4

H

- Hardware Redundancy
 - channel protection 3-7
 - controller cache memory 3-5
- heterogeneous hosts, overview and example 4-17
- host bus adapter utility 4-12
- host group, defined as a topological element 4-12
- host port, defined as a topological element 4-12
- host type, setting 4-17
- host, defined as a topological element 4-12

hot spare drive 3-6
hot-add utility 4-9

I

increase volume capacity, also called Dynamic
Volume Expansion (DVE) 4-10
I/O (Input/Output)
defined 3-4
multi-path driver w/AVT 3-8
password protection 3-9

L

Logical View
defined 2-11
status icon changes 5-12
status icons 5-3
LUN (logical unit number), defined 4-13

M

management domain 2-6
Mappings View, overview 4-11
media scan 3-6
medical imaging applications 3-4
menu options
Array Management Window 2-12
Enterprise Management Window 2-7
multimedia applications 3-4
multi-path driver w/AVT 3-8
multi-user environments 3-4

N

Needs Attention status 5-3

O

Optimal status 5-3
Overall Health Status 5-3

P

password protection 3-9
Physical View
defined 2-11
status icon changes 5-12

power outage 3-6
preferred controller owner 3-8
Problem Notification, indicators 5-8

R

RAID (Redundant Array of Independent Disks),
level description 3-3
reconstruction 3-6
Recovery Guru, overview and example 5-10
re-create snapshot volumes 4-7
Remote Volume Mirroring, overview 4-8

S

SANshare Storage Partitioning
overview and example 4-14
topology 4-12
Wizard 4-14
Script Editor 2-4
Snapshot Volumes
overview 4-6
Wizard 4-6
SNMP (Simple Network Management
Protocol), defined 2-4
Software Redundancy
Remote Volume Mirroring 3-2
Snapshot Volumes 3-2
standard volumes
Create Volume Wizard 4-4
overview 4-4
status icon, changes during recovery
procedure 5-12

T

toolbars
Array Management Window 2-13
Enterprise Management Window 2-8
topological elements, defined 4-11
Topology View 4-11

U

unconfigured capacity 4-5
Unresponsive status 5-3
UPS (Uninterruptible Power Supply) 3-5

V

volume

- creation overview 4-2
- definitions 4-2
- volume-to-LUN settings 4-16

volume group

- capacity of 3-3
- creation parameters 4-10
- defined 4-3

volume-to-LUN mapping, defined 4-13

W

write caching 3-5

WWN (World Wide Name) 4-12

