



IBM Netfinity Fibre Channel Storage Manager  
**Concepts Guide**





Netfinity Fibre Channel Storage Manager

**Concepts Guide**

Before using this information and product it supports, be sure to read the general information under "Appendix A. Notices".

**First Edition (December 1999)**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication was developed for products and services offered in the United States of America and the United Kingdom. It is possible that this publication may contain references to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Requests for technical information about IBM products should be made to your IBM reseller or IBM marketing representative.

No part of this publication may be reproduced or distributed in any form or by any means without prior permission in writing from the International Business Machines Corporation.

**©Copyright International Business Machines Corporation 1999. All rights reserved.**

Note to U.S. Government Users—Documentation related to restricted rights—Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

# Contents

---

<b>Figures</b> .....	<b>v</b>
<b>Tables</b> .....	<b>vii</b>
<b>Chapter 1. Understanding Storage Management Software</b> .....	<b>1</b>
Introducing Storage Subsystems .....	1
What is a Storage Subsystem? .....	1
Managing Storage Subsystems .....	4
Storage Subsystem Connectivity .....	4
Hardware Components .....	5
Software Components .....	5
About SM7client .....	6
Logical and Physical View Description .....	9
Network Management Connection Types .....	10
Populating a Management Domain .....	11
Configuring Storage Subsystems .....	11
Protecting Data .....	11
Creating Logical Drives and Storage Partitions .....	16
Example .....	19
Creating Storage Partitions .....	20
Monitoring Storage Subsystems .....	23
Failure Notification .....	25
Recovering from Storage Subsystem Problems .....	28
Recovery Guru Window .....	29

Tuning Storage Subsystems .....	32
Balancing the I/O Load .....	33
Optimizing the Transfer Rate .....	33
Optimizing the I/O Request Rate .....	34
Terminology Changes in This Release .....	37
<b>Chapter 2. Task Reference.....</b>	<b>39</b>
<b>Appendix A. Notices .....</b>	<b>47</b>

# Figures

---

Figure 1. Storage Subsystem Physical Components .....	2
Figure 2. Host-Agent Managed Storage Subsystem .....	4
Figure 3. Directly Managed Storage Subsystem .....	4
Figure 4. SM7client Windows .....	7
Figure 5. Enterprise Management Window .....	8
Figure 6. Subsystem Management Window Logical View and Physical View .....	9
Figure 7. Device Tree with a Management Domain .....	11
Figure 8. Storage Partition Configuration Example .....	19
Figure 9. Topology View .....	20
Figure 10. Define New Mapping Window .....	21
Figure 11. Creating Logical Drive-to-LUN Mapping in the Topology View .....	22
Figure 12. Mappings View .....	22
Figure 13. Storage Subsystem Status Icons .....	24
Figure 14. Failure Notification in the Logical View .....	26
Figure 15. Failure Notification in the Physical View .....	27
Figure 16. Recovering from Storage Subsystem Problems .....	28
Figure 17. Recovery Guru Window .....	29
Figure 18. Recovery Guru Window Showing Replaced Status .....	30
Figure 19. Recovery Guru Window Showing Optimal Status .....	31





# Tables

---

Table 1. Storage Subsystem Physical Components .....	2
Table 2. Storage Subsystem Logical Components .....	2
Table 3. Storage Management Architecture Hardware Components .....	5
Table 4. Major Software Components .....	5
Table 5. RAID Level Configurations .....	12
Table 6. Creating Logical Drives from Unconfigured or Free Capacity .....	16
Table 7. Storage Partitions Definitions Window Views .....	17
Table 8. Storage Partitions Terminology .....	17
Table 9. Storage Subsystem Status Icon Quick Reference .....	25
Table 10. Performance Monitor Data Quick Reference .....	32
Table 11. Terminology Change Reference .....	37
Table 12. Populating a Management Domain .....	39
Table 13. Configuring Storage Subsystems .....	39
Table 14. Using the Script Editor .....	40
Table 15. Configuring Storage Partitions .....	41
Table 16. Protecting Data .....	42
Table 17. Event Notification .....	43
Table 18. Recovering from Problems .....	44
Table 19. Miscellaneous System Administration .....	45

Table 20. Security .....	45
Table 21. Performance and Tuning .....	46

# Chapter 1. Understanding Storage Management Software

---

This chapter explains the concepts used in Netfinity<sup>®</sup> Fibre Channel Storage Manager 7.x and helps you understand this software.

Topics covered in this chapter include:

- "What is a Storage Subsystem?" on page 1
- "Managing Storage Subsystems" on page 4
- "Configuring Storage Subsystems" on page 11
- "Monitoring Storage Subsystems" on page 23
- "Recovering from Storage Subsystem Problems" on page 28
- "Tuning Storage Subsystems" on page 32
- "Terminology Changes in This Release" on page 37

"Chapter 2. Task Reference" on page 39 provides a task-based index to the appropriate online help systems.

## Introducing Storage Subsystems

This section describes the physical and logical components of a storage subsystem and introduces storage management tasks.

### What is a Storage Subsystem?

A *storage subsystem* is a storage entity that consists of a collection of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as arrays and logical drives). A storage subsystem can span multiple physical enclosures.

## Physical Components

Figure 1 shows the storage subsystem physical components.




Figure 1. Storage Subsystem Physical Components

Table 1 describes the storage subsystem physical components.


**Table 1: Storage Subsystem Physical Components**

Component	Description
<b>Drive</b>	An electro-magnetic mechanical device that provides the physical data storage media.
<b>Drive enclosure</b>	An enclosure that contains drives, power supplies, fans, environmental cards, and other supporting components.
<b>Controller</b>	A board and firmware that control the drives and implement the storage management functions.
<b>Controller enclosure</b>	An enclosure that contains one or more controllers, power supplies, fans, and other supporting components.




## Logical Components

The storage subsystem drives provide the physical storage capacity for data. Use the storage management software to configure the physical capacity into logical components that improve data availability and application performance. Table 2 describes the storage subsystem logical components.

**Table 2: Storage Subsystem Logical Components**

Component	Description
<b>Logical drive</b> 	A logical object, which is the basic structure you create on the storage subsystem to store data. A logical drive is a contiguous subsection of an array that is configured with a <i>RAID level</i> to meet application needs for data availability and I/O performance. A logical drive is seen by the operating system as one drive.

**Table 2: Storage Subsystem Logical Components (continued)**

<b>Array</b> 	A set of drives that are logically grouped together by the controllers in a storage subsystem to provide logical drives to host computers.
<b>Free capacity</b> 	A contiguous region of unassigned capacity on a designated array. The space is used to create one or more logical drives. <b>Note:</b> In the Subsystem Management window Logical View, free capacity is shown as Free Capacity Nodes. Multiple Free Capacity nodes can exist on an array.
<b>Unconfigured capacity</b> 	The capacity in the storage subsystem from drives that are not assigned to an array. The space is used to create new arrays and logical drives. <b>Note:</b> In the Subsystem Management window Logical View, unconfigured capacity is shown as an Unconfigured Capacity Node.

### What is Storage Subsystem Management?

Storage subsystem management activities include:

- Configuring available storage subsystem capacity into logical drives to maximize data availability, optimize application performance, and utilize storage resources.
- Granting access to partitions to the application hosts in the enterprise.
- Setting up a *management domain*, which is a collection of storage subsystems to manage.
- Monitoring storage subsystems in the management domain for problems or conditions that require attention.
- Configuring destinations to receive alert messages for critical problems concerning one or more storage subsystems in the management domain.
- Recovering from storage subsystem problems to maximize data availability.
- Tuning for optimal application performance.

# Managing Storage Subsystems

This section describes storage subsystem I/O and management connectivity and the hardware and software components in the storage management architecture. It also introduces Netfinity Fibre Channel Storage Manager 7.x and describes how storage subsystems are managed using this software.

## Storage Subsystem Connectivity

A storage subsystem receives data, or *I/O*, from the application host over Fibre Channel I/O data connections.




Figure 2. Host-Agent Managed Storage Subsystem

In previous versions of the storage management software, the software was installed on the application host. Storage management commands were sent from the host to the storage subsystem controllers over the I/O data connection.

Netfinity Fibre Channel Storage Manager 7.x is installed on personal computers or workstations, called *management stations*, anywhere on the corporate network. Storage management commands are sent from these management stations over an Ethernet network to the storage subsystem. The connection between the management station and the storage subsystem controller is called the *network management connection*. There are two kinds of network management connections, described in "Network Management Connection Types" on page 10.




Figure 3. Directly Managed Storage Subsystem

## Hardware Components

Table 3 describes the hardware components connected to the network in storage management architecture. For a detailed description of the network setup requirements, refer to the *Installation and Support Guide*.

**Table 3: Storage Management Architecture Hardware Components**

Component	Description
<b>Management station</b>	A computer used to manage storage subsystems on the network.
<b>Network management station (optional)</b>	<p>A network management station (NMS) is a console with installed SNMP-compliant network management software that receives and processes information about managed network devices using SNMP (Simple Network Management Protocol).</p> <p>The storage management software sends critical alerts (using <i>SNMP trap messages</i>) to configured destinations.</p>
<b>BOOTP or BOOTP-compliant DHCP server</b>	A BOOTP or BOOTP-compliant DHCP server is used to assign the network-specific information such as IP address and host name, for each controller. This server is not required if you are managing all storage subsystems through the I/O connection using a host-agent.
<b>Host</b>	A computer running one or more applications that accesses the storage subsystem over an I/O data connection.
<b>Storage subsystem</b>	A storage entity, managed by the storage management software, that consists of both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as arrays and logical drives).
<b>File server</b>	You can store the storage management software on a central file server. Then, management stations on the network can remotely access the storage management software.

## Software Components

Table 4 describes the three major software components of Netfinity Fibre Channel Storage Manager 7.x.

**Table 4: Major Software Components**

Component	Description
<b>SM7client</b>	<p>The SM7client software is installed on the management station, enabling the system administrator to perform storage management.</p> <p>For more information, see "About SM7client" on page 6.</p>

**Table 4: Major Software Components (continued)**

---

<b>SM7agent</b>	<p>The SM7agent, or host-agent management software, is a software component that you can install on one or more hosts connected to the storage subsystems. The host-agent, along with the network connection on the host, provides a network management connection to the storage subsystem instead of the individual Ethernet connections on each controller. The management station can communicate with a storage subsystem through the host that has host-agent management software installed. The host-agent receives requests from the management station through the network connection to the host and sends them to the controllers in the storage subsystem through the fibre channel I/O connection.</p> <p>The SM7agent package contains the SM7devices utility that correlates the logical drives you create using the storage management software with their operating system device names. Always install the SM7agent software on all host computers, even if you plan to manage storage subsystem directly over the network. You can stop the SM7agent from running using an operating system-specific method. For more information, refer to the <i>Installation and Support Guide</i>.</p> <p><b>Note:</b> Hosts with the host-agent software installed are automatically discovered by the storage management software and appear in the device tree in the Enterprise Management window, along with their attached storage subsystems. A storage subsystem might be duplicated in the device tree if you are managing it through its Ethernet connections and it is attached to a host with the host-agent software installed. In this case, the duplicate storage subsystem icon can be removed from the device tree using the Remove Device option in the Enterprise Management window.</p>
<b>Redundant Dual Active Controller (RDAC) multi-path driver software</b>	<p>RDAC is an I/O path failover driver installed on the host computers that access the storage subsystem. Usually, a pair of active controllers is located in a storage subsystem. Each logical drive in the storage subsystem is “owned” by one controller that controls the I/O between the logical drive and the application host along the I/O path. When a component in the I/O path fails, such as a cable or the controller itself, the RDAC multi-path driver transfers ownership of the logical drives assigned to that controller to the other controller in the pair.</p> <p><b>Note:</b> Some Operating Systems have built-in I/O Data connection failover drivers and do not require this multi-path driver.</p>

---

## About SM7client

SM7client is called *thin* because it provides an interface for storage management based on information supplied by the storage subsystem controllers. When you use SM7client from a management station to manage a storage subsystem, you send commands to the storage subsystem controllers. The controller firmware contains the necessary logic to carry out the storage management commands. The controller is responsible for validating and executing the commands and providing the status and configuration information that is sent back to SM7client.



SM7client has two main windows: the Enterprise Management window and the Subsystem Management window, as shown in Figure 4.

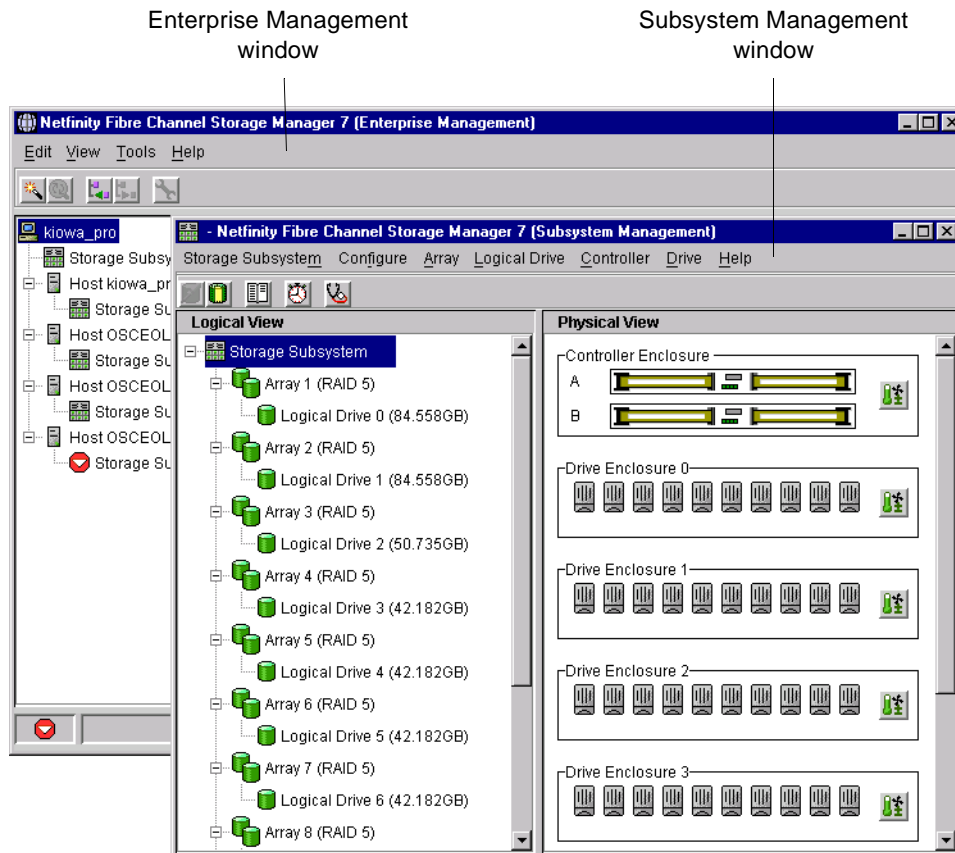


Figure 4. SM7client Windows

## Enterprise Management Window

The Enterprise Management window is the first window to appear when you start the storage management software. Use the Enterprise Management window to:

- Add and discover the storage subsystems you want to manage.
- Provide a comprehensive view of all storage subsystems in the management domain.
- Perform batch storage subsystem management tasks using the Script Editor.

A local configuration file named `emwdata.bin` stores a list of the storage subsystems to include in the management domain. After adding the storage subsystems, use the Enterprise Management window primarily for coarse-level monitoring and alert notification of non-optimal storage subsystems. Also, it is the home base for launching a Subsystem Management window for a particular storage subsystem. You can perform all subsystem-specific management tasks using the Subsystem Management window.

Figure 5 shows the Enterprise Management window.

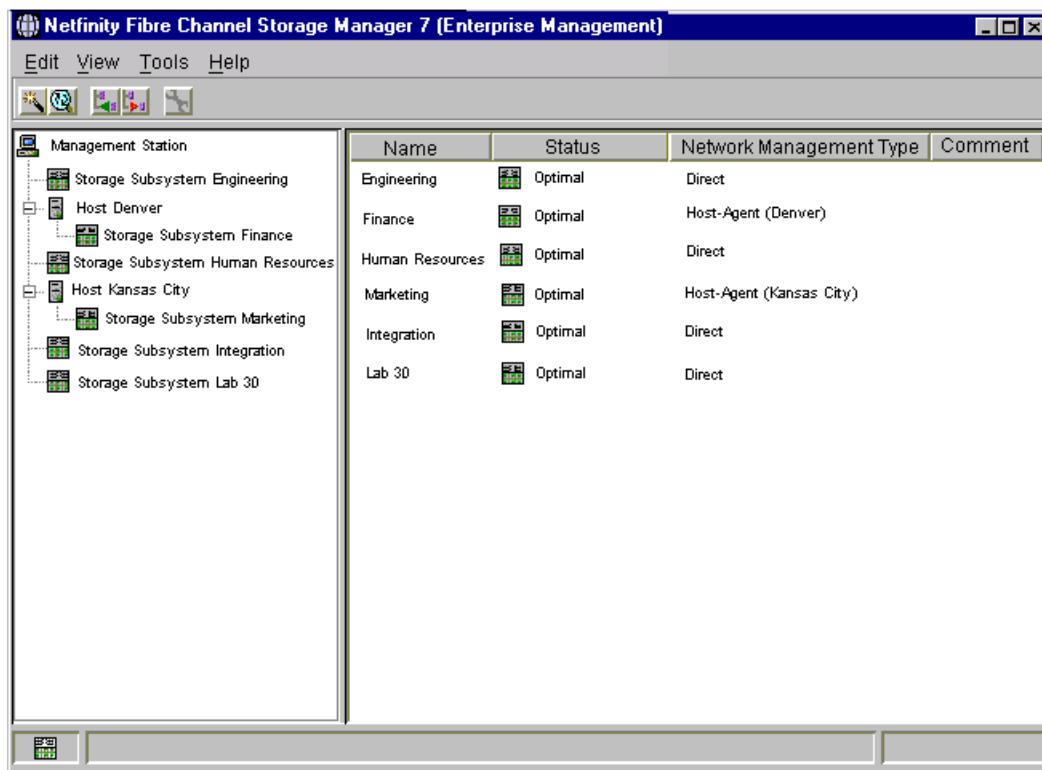


Figure 5. Enterprise Management Window

## Subsystem Management Window

Use the Subsystem Management window to manage (configure, monitor, recover, and tune) the logical and physical components of a storage subsystem. The Subsystem Management window is specific to an individual storage subsystem. Therefore, you can manage only a single storage subsystem within a Subsystem Management window. You can launch other Subsystem Management windows to manage other storage subsystems.

Figure 6 shows the Logical View and Physical View of the Subsystem Management window.

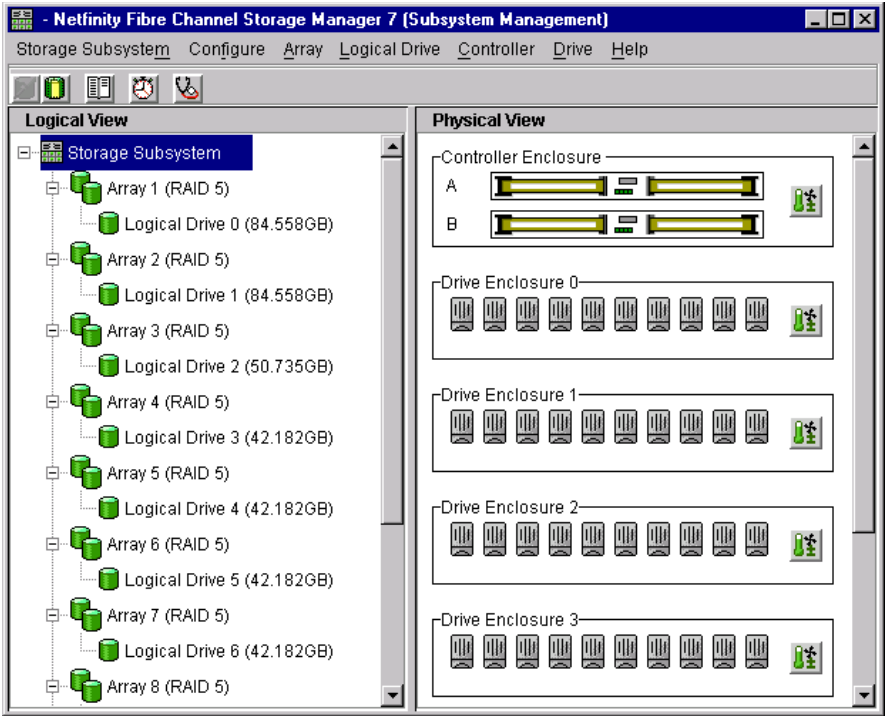


Figure 6. Subsystem Management Window Logical View and Physical View

### Logical and Physical View Description

The *Logical View* displays in the left pane of the Subsystem Management window and shows the storage subsystem capacity organized into arrays and logical drives. The *Physical View* displays in the right pane of the Subsystem Management window and shows the physical devices in the storage subsystem, such as controllers, drives, and other components. Selection of a logical drive or other entity in the Logical View shows you the associated physical components in the Physical View.

## Network Management Connection Types

The storage management software manages all storage subsystems over the network. The software allows two different types of network management connections to a storage subsystem: direct or host-agent.

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

### Directly Managed Storage Subsystems

In this management method, the controllers in the storage subsystem are managed directly over the network through each controller's Ethernet connection. To manage the storage subsystem through these Ethernet connections, you must define each controller's IP address (or optionally, each controller's host name) and cable the Ethernet connections on the storage subsystem to the network. Then, use the Enterprise Management software to add the storage subsystem to the management domain.

### Host-Agent Managed Storage Subsystems

In this management method, the controllers in the storage subsystem are managed through a network connection to a host, rather than through the Ethernet connections to each controller. The host-agent software on the host facilitates communication between the management station and the controllers in the storage subsystem. To manage a storage subsystem using this method, you must install the host-agent software on the host and then use the Enterprise Management window to add the host to the management domain. By including the host in the domain, you also include attached host-agent managed storage subsystems.

## Populating a Management Domain

A management domain is a collection of storage subsystems that you want to manage. The Enterprise Management window device tree displays the management domain, as shown in Figure 7.




Figure 7. Device Tree with a Management Domain

There are two ways to populate a management domain:

- Use the **Automatic Discovery** option to automatically discover directly managed and host-agent managed storage subsystems on the local subnetwork and add them to the management domain. The Enterprise Management software discovers host-agent managed storage subsystems by first discovering the hosts that provide network management connections to the storage subsystems. Then the host and associated storage subsystems display in the device tree.
- Use the **Add Device** option. For a directly managed storage subsystem, enter a host name or IP address for each controller in the storage subsystem. For a host-agent managed storage subsystem, enter a name or IP address for the host that is attached to the storage subsystem.

For more information about populating a management domain, refer to the Enterprise Management window online help system.

## Configuring Storage Subsystems

This section describes the storage subsystem configuration options that are used to maximize data availability. It also outlines the high-level steps to configure available storage subsystem capacity into logical drives and storage partitions.

## Protecting Data

Storage subsystems are designed for reliability and feature redundant drives, controllers, power supplies, and fans. This hardware redundancy keeps the storage subsystem working if a component fails. In addition, the storage management software features several options that you can use to protect data.

The following strategies help maximize data protection and availability:

- Using appropriate RAID levels for data redundancy
- Protecting data in the controller cache memory
- Configuring hot spare drives
- Enabling a background media scan
- Protecting the I/O data paths
- Configuring storage subsystem password protection
- Configuring alert destinations for storage subsystem critical event notification
- Using available channel protection

These strategies are described in the following sections.

### Using Appropriate RAID Levels for Data Redundancy

RAID levels 1, 3, and 5 write redundancy data to the drive media for fault tolerance. The redundancy data might be a copy of the data (mirrored) or an error correcting code derived from the data. The redundancy data is stored on a different drive from the user data that it protects. It is used to quickly reconstruct information on a replacement drive, if a drive fails.

RAID relies on a series of configurations, called Levels, to determine how user and redundancy data is written and retrieved from the drives. This storage management software offers four RAID level configurations: RAID level 0, 1, 3, and 5. Each level provides different performance and protection features.

Table 5 describes the RAID level configurations available with Netfinity Fibre Channel Storage Manager 7.x.

**Table 5: RAID Level Configurations**

<b>RAID Level</b>	<b>Description</b>
RAID 0	RAID level 0 offers simplicity and high performance, but does not provide data redundancy. A Level 0 array spreads data across all drives in the array. If one drive in the array fails, all logical drives contained on the array fail.
RAID 1	RAID level 1 uses disk mirroring to make an exact copy of data from one drive to another drive. If one drive fails in a RAID 1 array, the mirrored drive takes over.  This RAID level is costly in terms of capacity but features outstanding performance.

**Table 5: RAID Level Configurations (continued)**

---

RAID 3	<p>RAID level 3 requires one dedicated disk in the logical drive to hold redundancy information. User data is striped across the remaining drives.</p> <p>RAID level 3 is a good choice for applications such as multimedia or medical imaging that write and read large sequential chunks of data. Usually the I/O size is large and drives operate in parallel to service a single request, delivering high I/O transfer rates.</p>
RAID 5	<p>RAID level 5 spreads redundancy information across all of the drives in the logical drive.</p> <p>RAID level 5 is a good choice in multi-user environments such as database or filesystem storage, where the I/O size is small and there is a high proportion of read activity. When the I/O size is small and the segment size is appropriately chosen, a single read request is retrieved from a single individual drive, leaving other drives available to concurrently service other I/O read requests and deliver fast read I/O request rates.</p>

---

**Note:** Within one array, a single RAID level is used and all redundancy data for that array is stored within the array.

The capacity of the array is the aggregate capacity of the member drives, minus the capacity reserved for redundancy data. The amount of capacity needed for redundancy depends on the RAID level used.

### **Protecting Data in the Controller Cache Memory**

*Write caching* allows the controller's cache memory to store write operations from the host, which increases system performance. However, a controller can fail with user data in its cache that has not been transferred to media. Also, the cache memory can fail while it contains unwritten data. *Write cache mirroring* protects the system from either of these possibilities. Write cache mirroring allows cached data to be mirrored across two redundant controllers with the same cache size. The data written to the cache memory of one controller is also written to the cache memory of the other controller. Therefore, if one controller fails, the other can complete all outstanding write operations. You can enable the **Write Cache Mirroring** parameter for each logical drive.

To prevent data loss or corruption, the controller writes cache data to disk (flushes the cache) periodically. When the cache holds a specified start percentage of unwritten data, a flush is triggered. When the cache flushes down to a specified stop percentage, the flush is stopped. For example, you can specify that the controller start flushing the cache when the cache reaches 80% full and stop flushing the cache when the cache reaches 16% full. For maximum data safety, you can choose low start and stop percentages. However, in both cases this increases the chance that data needed for a host read will not be in the cache, decreasing the cache hit percentage and therefore the I/O request rate. It also increases the number of disk writes necessary to maintain the cache level, increasing system overhead and further decreasing performance.

If a power outage occurs, data in the cache that is not written to the drive media is lost, even if it is mirrored to the cache memory of both controllers. For protection against this occurrence, there are batteries in the controller enclosure that protect the cache against power outages. Change the controller enclosure batteries at the recommended time intervals. The storage management software features a battery age clock you can set when you replace a battery. This clock keeps track of the age (in days) of the battery so you know when it is time to replace the battery.

Sometimes write caching is disabled when batteries are low or discharged. If you enable a parameter called **write caching without batteries** on a logical drive, write caching continues even when batteries in the controller enclosure are discharged.

---

**Important:** For maximum data availability, do not enable this parameter because data in the cache is lost during a power outage if the controller enclosure does not have working batteries.

---

## Configuring Hot Spare Drives

You can replace a failed hard disk drive while the storage subsystem is operating. If the drive is part of a RAID 1, 3, or 5 array, the controller uses redundancy data to automatically reconstruct the data from the failed drive to the replacement drive. This is called *reconstruction*.

You can assign available drives in the storage subsystem as *hot spare* drives to keep data available. A hot spare is a drive containing no data that acts as a standby in case a drive fails in a RAID 1, 3, or 5 logical drive. The hot spare drive adds another level of redundancy to the storage subsystem. If a drive fails in the storage subsystem, the hot spare is automatically substituted for the failed drive without requiring a physical swap. If the hot spare drive is available when a drive fails, the controller uses redundancy data to reconstruct the data from the failed drive to the hot spare drive. When you have physically replaced the failed drive, the data from the hot spare is copied back to the replacement drive. This is called *copyback*.

## Enabling a Background Media Scan

A media scan is a background process that runs on all logical drives in the storage subsystem for which it is enabled, providing error detection on the drive media. The advantage of enabling the media scan process is that the process can find media errors before they disrupt normal drive reads and writes. The media scan process scans all logical drive data to verify that it is accessed, and optionally scans the logical drive redundancy data.



## **Protecting the I/O Data Path**

If the operating system on the application host does not have a native multi-path failover driver, use the RDAC multi-path driver that is provided with the storage management software.

The RDAC multi-path driver manages the I/O data path failover process for storage subsystems with redundant controllers. If a component (for example, a cable, controller, host adapter, and so on) fails along the I/O data path, the RDAC multi-path driver automatically reroutes all I/O operations to the other controller.

## **Using Available Channel Protection**

When you create logical drives from unconfigured capacity, array candidates are shown in the **Create Logical Drive - Select Array** dialog box with information about whether or not the array candidate has channel protection. In a SCSI environment, channel protection depends on the RAID level of the array and how many drives are present on any single drive channel. For example, a RAID 5 array does not have channel protection if more than one drive is present on a single drive channel.

In a Fibre Channel environment, an array candidate has channel protection because there are redundant Fibre Channel Arbitrated Loops when the storage subsystem is properly cabled.

## **Configuring Storage Subsystem Password Protection**

You can configure a storage subsystem password that must be supplied when a destructive operation is attempted. Use a long password to increase security. Be sure to use a minimum of 15 alphanumeric characters.

For additional security, the storage management software encrypts the password using a secure hash algorithm.

## **Configuring Alert Destinations for Storage Subsystem Critical Event Notification**

There are flexible options available for configuring alert notification destinations. You can set up alert notification destination addresses for notifications about:

- Storage subsystems in the management domain
- Storage subsystems attached and managed through a particular host
- Individual storage subsystems

The storage management software also enables you to validate potential destination addresses and specify management domain global e-mail alert settings for mail server and sender e-mail address.

## Creating Logical Drives and Storage Partitions

You can configure logical drives from either unconfigured capacity or free capacity in a storage subsystem, as described in Table 6.

**Table 6: Creating Logical Drives from Unconfigured or Free Capacity**

<b>If you create the logical drive from...</b>	<b>Then...</b>
<b>Unconfigured capacity</b>	Do the following: <ol style="list-style-type: none"><li>1. Define the array, specifying two parameters: RAID level and capacity (how large you want the array, that is, how many drives you want to include in the array).</li><li>2. Specify the logical drive parameters (logical drive name, usage, desired capacity, controller ownership, and storage partition mapping preference).</li></ol>
<b>Free capacity</b>	Because free capacity only resides on an existing array, the overall capacity (that is, the number of drives) and RAID level are already specified.  Therefore, the only step that is required is to specify the parameters (logical drive name, usage, desired capacity, controller ownership, and storage partition mapping preference).

### Creating Storage Partitions

With the storage partitions feature of the storage management software, you can consolidate logical drives into sets of logical drives called *storage partitions*. You grant visibility of partitions to defined hosts or a defined set of hosts called a *host group*. Storage partitions enable hosts to share storage capacity, consolidating storage and reducing storage management costs.

The Storage Partitions Definitions window is used to create storage partitions. Table 7 describes the two views of the Storage Partitions Definitions window.

**Table 7: Storage Partitions Definitions Window Views**

<b>View</b>	<b>Description</b>
<b>Topology</b>	<p>The Topology View shows defined topological elements (host ports, hosts, host groups) and the default host group. To use storage partitions you must define topological elements.</p> <p>If you are not using storage partitions, you can use the Topology View to define the hosts and use the Mappings View to show the default LUN numbers used to access logical drives in the storage subsystem.</p>
<b>Mappings</b>	<p>The default and specific logical drive-to-LUN mappings in the storage subsystem display in table format in the Mappings View. Information displayed about the logical drives includes: the topological entity that can access the logical drive, logical drive name, unique logical drive identifier, and LUN number used to access the logical drive.</p>

Table 8 describes the terminology that is used in the Storage Partitions Definitions window.

**Table 8: Storage Partitions Terminology**

<b>Term</b>	<b>Description</b>
<b>Storage partition</b>	<p>A collection of storage subsystem logical drives that is visible to a host or is shared among hosts that are part of a host group.</p>
<b>Storage partition topology</b>	<p>A collection of default host group, host group, host and host port nodes in the Topology View of the Storage Partitions window.</p> <p>You must define the host port, host, and host group topological elements in order to grant access to hosts and host groups using logical drive-to-LUN mappings.</p>
<b>Host port</b>	<p>Host ports physically reside on the host adapters and are automatically discovered by the storage management software. To give a host access to a partition, you <i>must</i> define its associated host ports. The host ports request data from a logical drive on behalf of the host computer; therefore, without associated host ports, a host cannot be given a logical drive-to-LUN mapping or request data from a logical drive using a LUN.</p> <p>Initially, all discovered host ports belong to the default host group.</p>

**Table 8: Storage Partitions Terminology (continued)**

<b>Host</b>	<p>An entity in the Storage Partitions topology that is a defined, logical set of host ports that need to access a logical drive.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"><li>• A defined host usually corresponds to a single computer running one or more applications that accesses a storage subsystem.</li><li>• A host does not need to belong to a defined host group unless the host needs to share access to a partition with other hosts.</li></ul>
<b>Host group</b>	<p>An entity in the Storage Partitions topology that defines a logical collection of hosts that need shared access to one or more logical drives.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"><li>• You can define a host group that corresponds to a <i>cluster</i>, or set of hosts that provide failover capabilities for each other.</li><li>• Hosts in a defined host group can be granted access to partitions independently of the host group. (Logical drive-to-LUN mappings are made to the host group or to an individual host in a host group.)</li></ul>
<b>Default host group</b>	<p>A logical collection of discovered host ports, defined hosts, and defined host groups in the Storage Partitions topology that:</p> <ul style="list-style-type: none"><li>• Are not involved in specific logical drive-to-LUN mappings</li><li>• Share access to logical drives with default logical drive-to-LUN mappings</li></ul>
<b>LUN</b>	<p>Logical Unit Number. The number a host uses to access a logical drive. Each host has its own LUN address space.</p>
<b>Logical drive-to-LUN mapping</b>	<p>An association of a logical drive with a single LUN. When you create a logical drive-to-LUN mapping, you specify both the LUN number that is used to access the logical drive and the defined host or host group that can access the logical drive.</p>
<b>Default logical drive-to-LUN mapping</b>	<p>This mapping allows host groups or hosts that do not have specific logical drive-to-LUN mappings (hosts or host groups that belong to the default host group) to access this logical drive.</p> <p>Logical drives are given default logical drive-to-LUN mappings when created. Legacy logical drives (created using previous versions of the storage management software) are automatically given default logical drive-to-LUN mappings but are given specific logical drive-to-LUN mappings with this version of the storage management software.</p>
<b>Storage partitions mapping preference</b>	<p>You can choose one of the following storage partition mapping preferences when creating a logical drive:</p> <ul style="list-style-type: none"><li>• Default logical drive-to-LUN mapping (described previously)</li><li>• No mapping. Choose this option when you create storage partitions to define a specific logical drive-to-LUN mapping for this logical drive.</li></ul>

## Example

Figure 8 shows three storage partition configurations. The first partition consists of logical drive *Financial* and is accessed by host KC-B. The second partition consists of logical drives *Legal* and *Engineering* and is accessed by the hosts in Host Group Kansas City. The third partition consists of logical drives *Marketing* and *HResources* and is accessed by the hosts in Host Group Omaha.

For information about how to configure the following example, see "Creating Storage Partitions" on page 20.

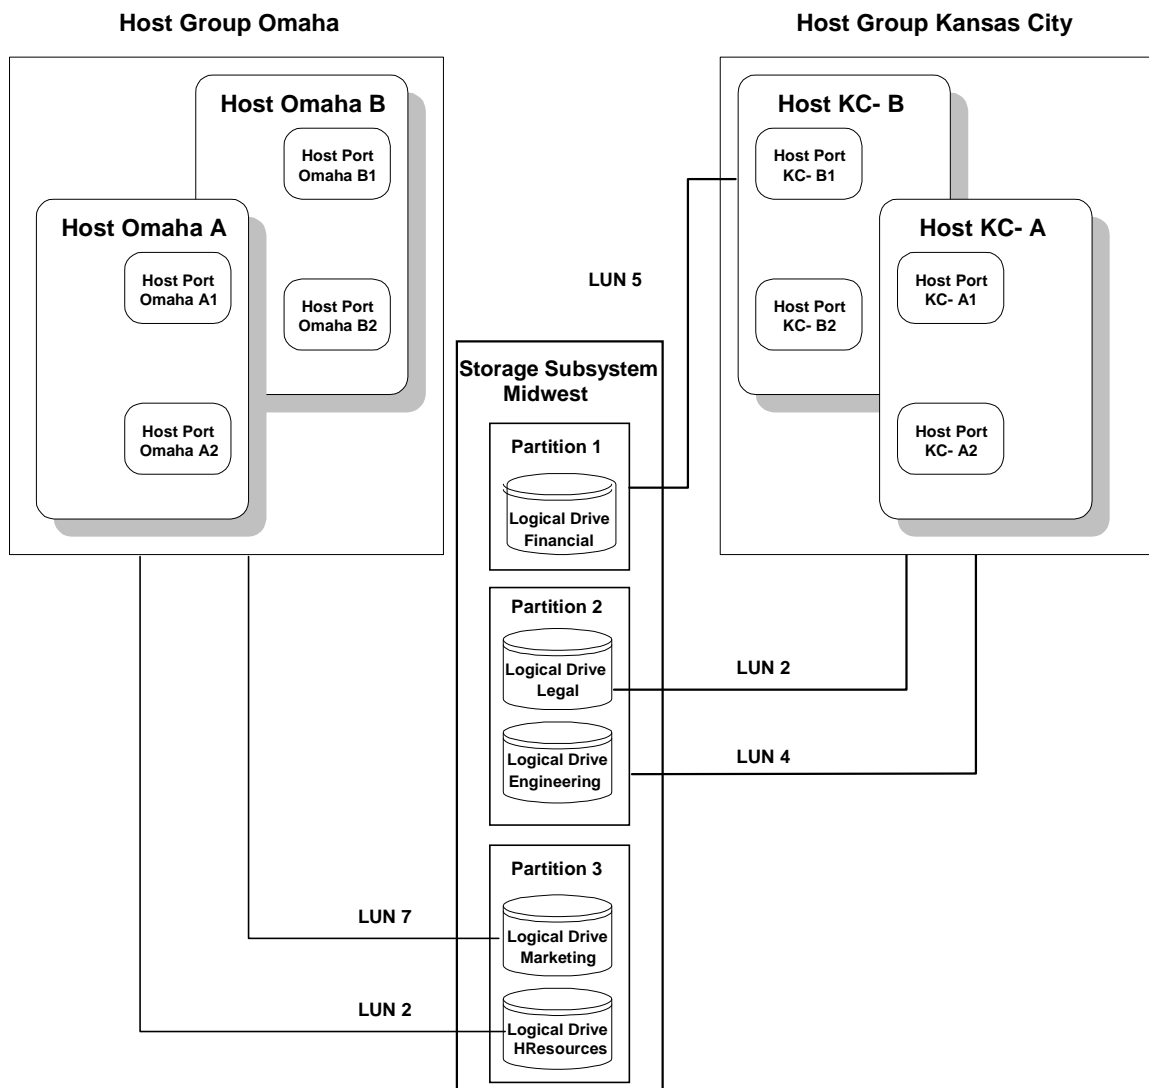


Figure 8. Storage Partition Configuration Example

## Creating Storage Partitions

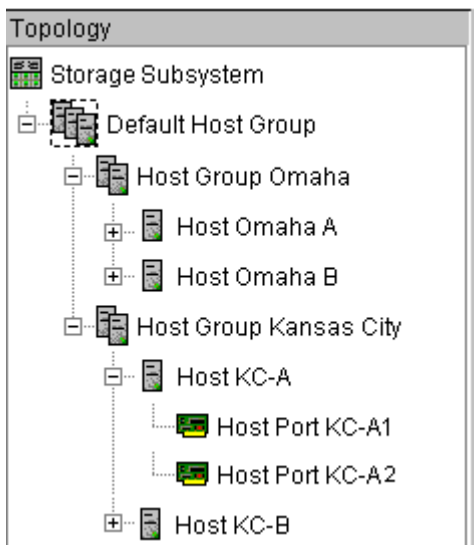
To create storage partitions, you must perform the following tasks:

1. Define storage partitions topology.
2. Define logical drive-to-LUN mappings.

### Defining Storage Partitions Topology

Use the Storage Partitions Definitions window to define the host groups, hosts, and associated host ports for which you want to define specific logical drive-to-LUN mappings (that is, the storage partitions *topology*).

The topological elements that you define (for example, host groups, hosts, and host ports) and the default host group are displayed in the Topology View, as shown in Figure 9.



In this window, the Host Groups Omaha and Kansas City and the hosts Omaha A, Omaha B, KC-A, and KC-B are defined in the Topology View.

The Host Ports KC-A1 and A2 (shown) and KC-B1, KC-B2, Omaha A-1, Omaha A-2, Omaha B-1, and Omaha B-2 (not shown) are also defined.

**Note:** Defined topological elements appear under the Default Host Group, because no specific logical drive-to-LUN mappings are defined.

Figure 9. Topology View

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

- Logical drives to which access is granted through a specific logical drive-to-LUN mapping
- Logical drives with default logical drive-to-LUN mappings. The host group or host is part of the default host group.

In Figure 9, the Host Groups Omaha and Kansas City and their associated hosts and host ports are not given specific logical drive-to-LUN mappings, so they appear under (as part of) the default host group and can access any logical drives that have default logical drive-to-LUN mappings. If the defined host groups and hosts are given specific logical drive-to-LUN mappings, they do not appear under the default host group, and they are not able to access logical drives with default logical drive-to-LUN mappings. (A host or a host group either can access logical drives with default logical drive-to-LUN mappings, or specific logical drive-to-LUN mappings.)

**Defining Logical Drive-to-LUN Mappings**

After you define the topology, the next step in creating storage partitions is to grant access to one or more logical drives to a defined host or a host group. You do this by defining specific logical drive-to-LUN mapping according to the following steps:

1. Select a defined host or a host group in the Topology View.
2. Access the **Define New Mappings** option.
3. Select a logical drive that you want the host group or host to access.
4. Choose the logical unit number (LUN) you want the host group or host to use to access the logical drive.

**Note:** Each host has its own LUN address space, so you can use the same LUN in more than one logical drive-to-LUN mapping, as long as that LUN is available for use by each particular host participating in the mappings.

Figure 10 shows the Define New Mapping dialog box.

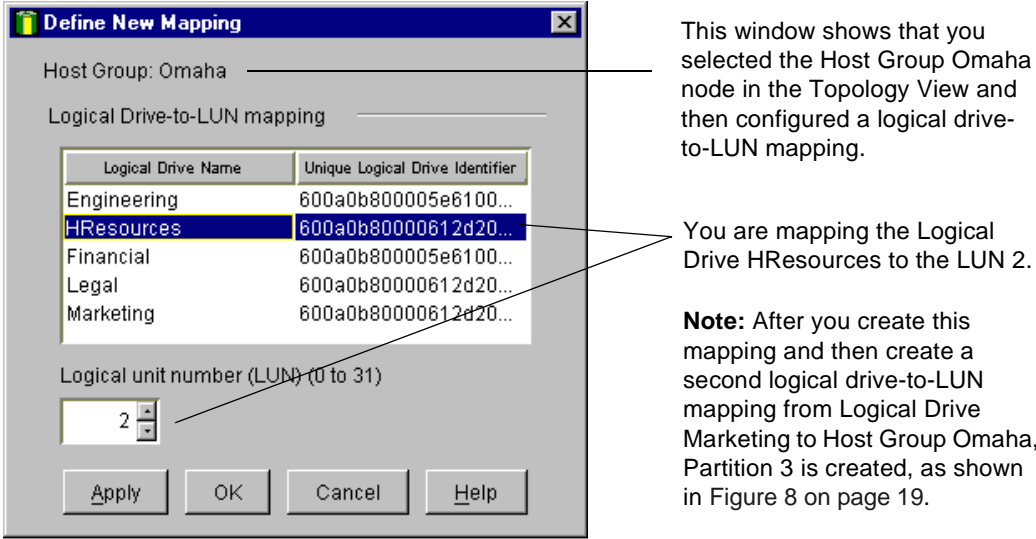


Figure 10. Define New Mapping Window

## Topology View Results

Figure 11 shows the results of creating logical drive-to-LUN mappings in the Topology View.

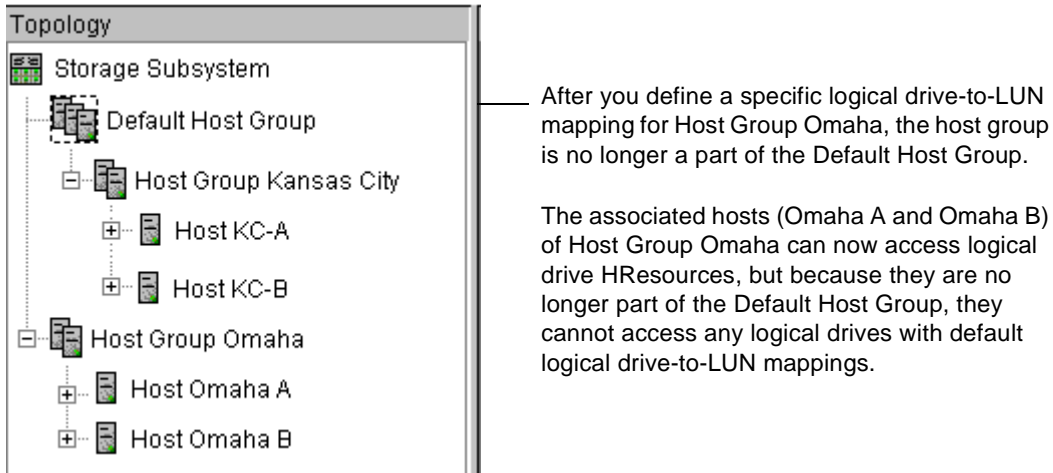


Figure 11. Creating Logical Drive-to-LUN Mapping in the Topology View

## Mapping View Results

After selecting the Host Group Omaha in the Topology View, the results display in the Mappings View, as shown Figure 12.

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

Figure 12. Mappings View

For more information about creating logical drives or storage partitions, see the Subsystem Management window online Help system.



## Monitoring Storage Subsystems

This section describes the monitoring of storage subsystems in a management domain.

Use the Enterprise Management window to:

- Monitor the coarse-level status of each individual storage subsystem and the overall health of all of the storage subsystems in the management domain
- Configure alert destinations for storage subsystem critical event notification

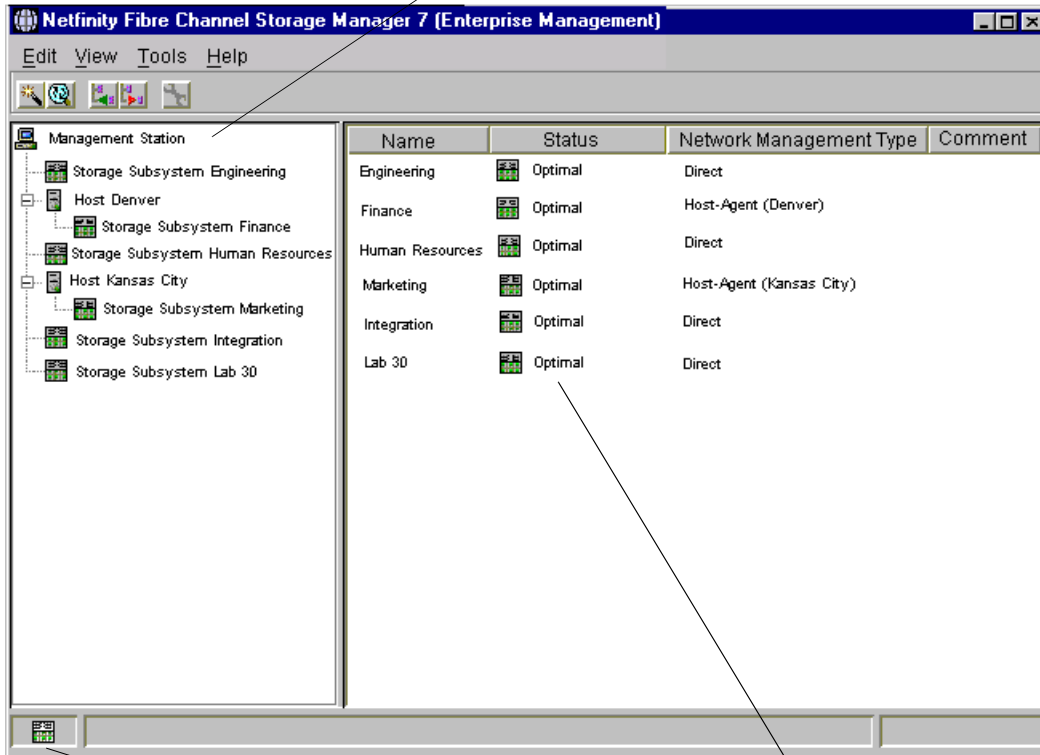
Use the Subsystem Management window to:

- Monitor the individual logical and physical components in an individual storage subsystem
- Monitor storage subsystem performance. For more information, see "Tuning Storage Subsystems" on page 32.

## Monitoring the Health of Storage Subsystems in the Management Domain

Figure 13 describes the storage subsystem status icons that display in the Enterprise Management window.

The Device Tree provides storage subsystem status for every storage subsystem in the management domain. (A status is shown for each network management connection.)



The Overall Health Status is a consolidated health status for all of the known storage subsystems in the management domain.

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






Figure 13. Storage Subsystem Status Icons

## Storage Subsystem Status Quick Reference

Table 9 on page 25 describes the storage subsystem status icons that display in the following:

- Device Tree, Device Table, and Overall Health Status area of the Enterprise Management window
- Root Node of the Logical View tree in the Subsystem Management window

**Table 9: Storage Subsystem Status Icon Quick Reference**

Status	Icon	Description
<b>Optimal</b>		Indicates that each component in the storage subsystem is in the desired working state.
<b>Needs Attention</b>		Indicates a problem on a storage subsystem that requires intervention. Launch the Subsystem Management window for the particular storage subsystem and then use the Recovery Guru to identify the cause of the problem and obtain appropriate instructions.
<b>Fixing</b>		Indicates that a Needs Attention condition is corrected and the storage subsystem is currently transitioning to an Optimal state (for example, a reconstruction operation is in progress). A Fixing status requires no action unless you want to check the progress of the operation in the Subsystem Management window. <b>Note:</b> Some recovery actions cause the storage subsystem state to change directly from Needs Attention to Optimal, without an interim state of Fixing.
<b>Unresponsive</b>		Indicates that the management station cannot communicate with the only controller or both controllers over its network management connection to the storage subsystem.
<b>Contacting Device</b>		Indicates that you have started the Enterprise Management window and the management software is establishing contact with the storage subsystem.

## Failure Notification

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

- A Needs Attention storage subsystem icon displayed in:
  - The Overall Health Status area, Device Tree View, or Device Table of the Enterprise Management window
  - The Subsystem Management window Logical View

- The Recovery Guru toolbar button in the Subsystem Management window changes from Optimal to Needs Attention status and flashes.
- Non-optimal component icons are displayed in the Subsystem Management window Logical View and Physical View
- You receive critical SNMP or e-mail error messages
- The hardware displays fault lights

### Example

Figure 14 shows the Subsystem Management window Logical View and the failure icons for Storage Subsystem Engineering.

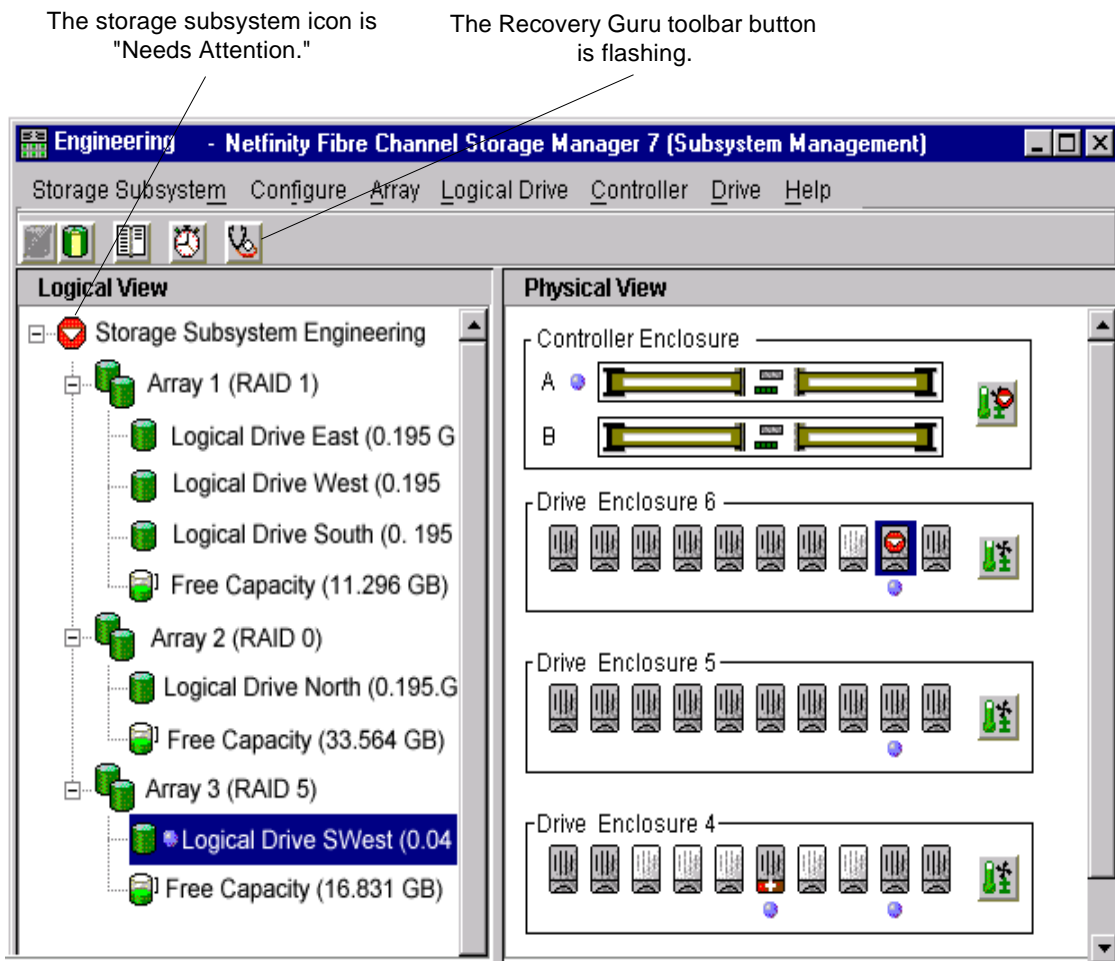
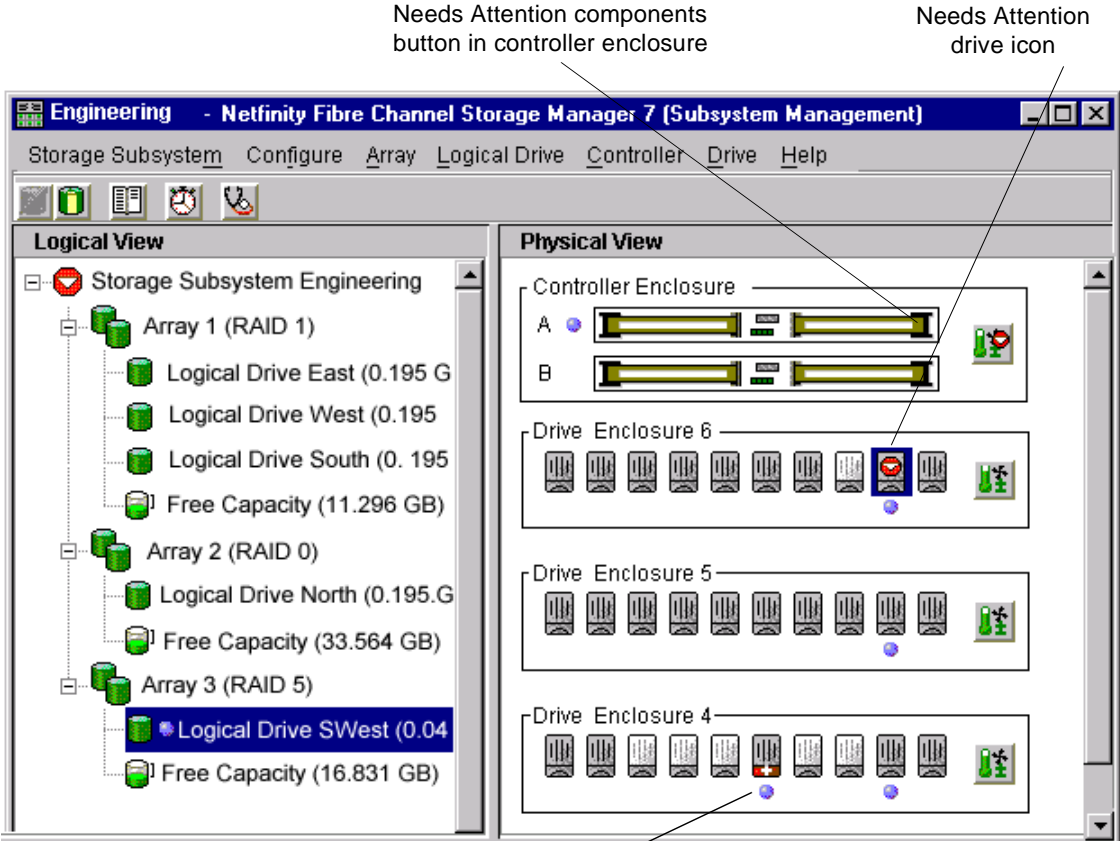


Figure 14. Failure Notification in the Logical View

Figure 15 shows the Subsystem Management window Physical View and the failure icons for Storage Subsystem Engineering.



This indicates that the hot spare drive is In Use (taking over for a failed drive), which means that the data from the failed drive has been reconstructed on the hot spare drive.

Figure 15. Failure Notification in the Physical View

You might receive failure notifications about this storage subsystem at the network management station or in e-mail. Also, hardware fault lights display on the affected controller and drive enclosures.

# Recovering from Storage Subsystem Problems

After you are notified of a problem, use the Recovery Guru to return the storage subsystem to an Optimal State. Figure 16 shows how to display the Recovery Guru window.

To display the Recovery Guru window, select **Recovery Guru** from the toolbar, or the **Storage Subsystem >> Recovery Guru** menu option.

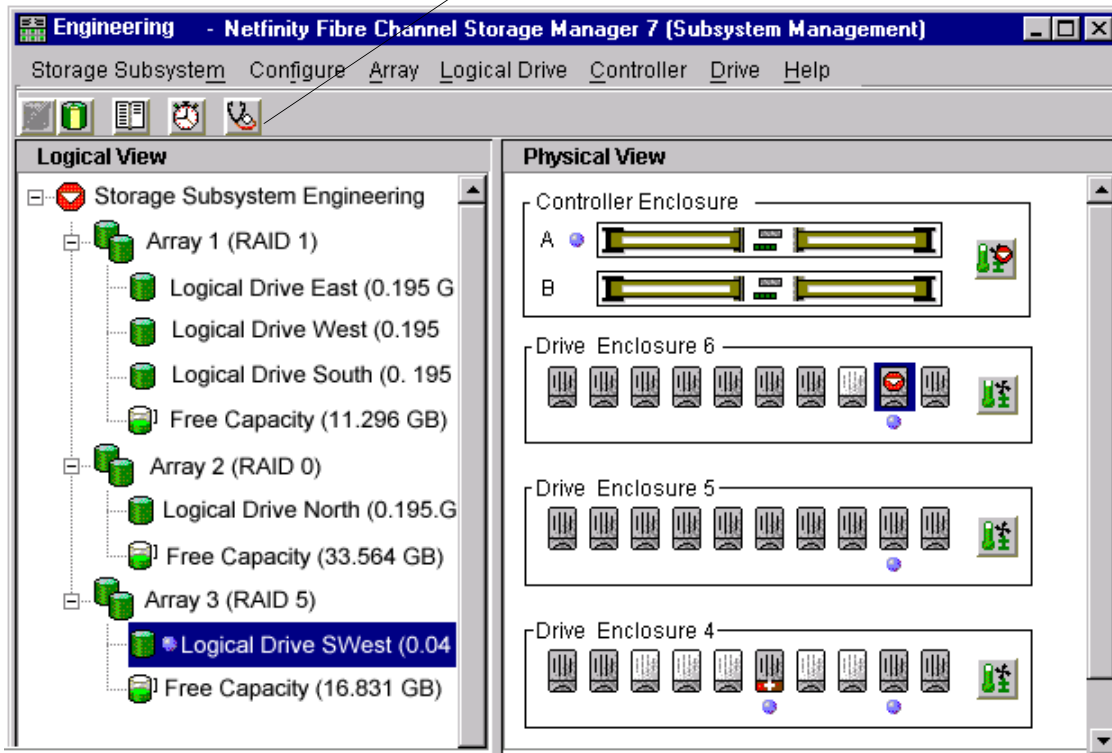


Figure 16. Recovering from Storage Subsystem Problems

## Recovery Guru Window

Figure 17 shows the Recovery Guru window.

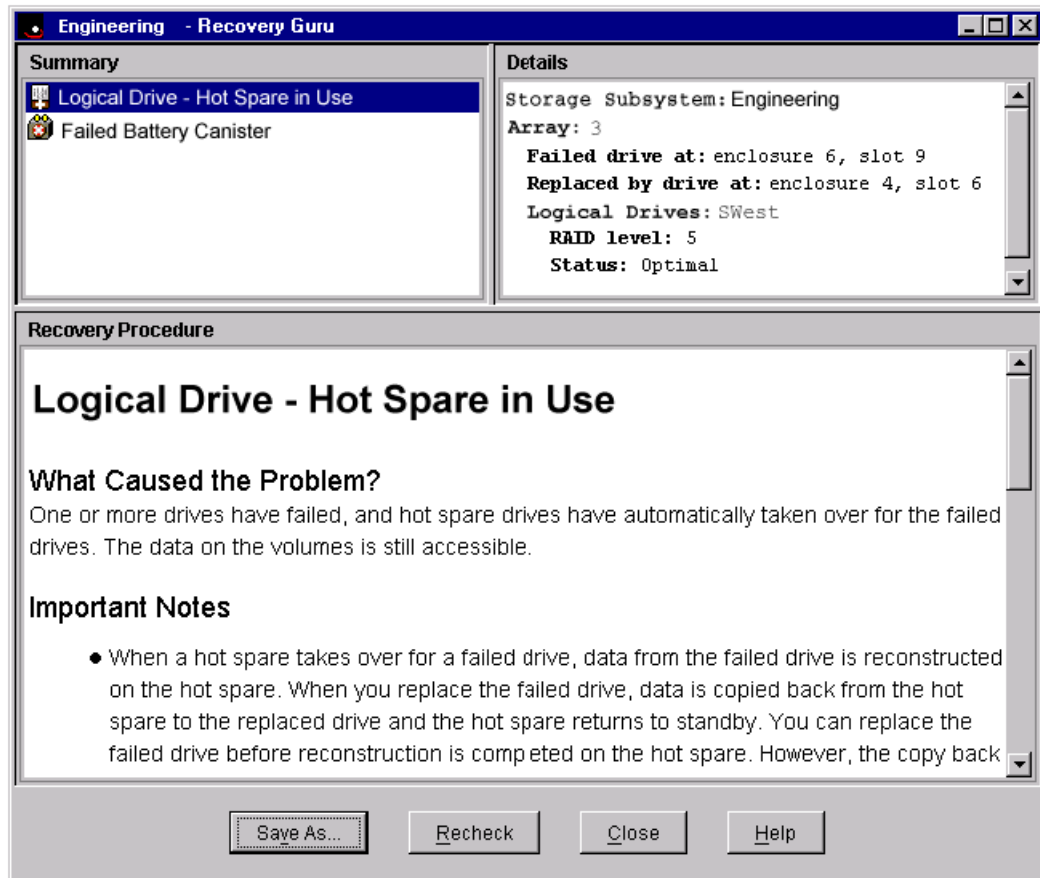


Figure 17. Recovery Guru Window

The Summary Area shows that there are two different failures in this storage subsystem. First, in Logical Drive SWest there is a hot spare drive taking over for a failed drive in enclosure 6, slot 9. Second, there is a failed battery canister in the controller enclosure. (If you select that failure in the Summary area, Details and a Recovery Procedure are shown.)

For each failure in the Summary area, a Recovery Procedure is given. Follow the steps in the Recovery Procedure to recover from each individual failure.

As you follow the Recovery Procedure to replace the failed drive, the associated logical drive (SWest) changes to Operation in Progress and the replaced drive changes to Replaced Drive status. The data that was reconstructed to the Hot Spare drive is copied back to the replaced drive, as shown in Figure 18.

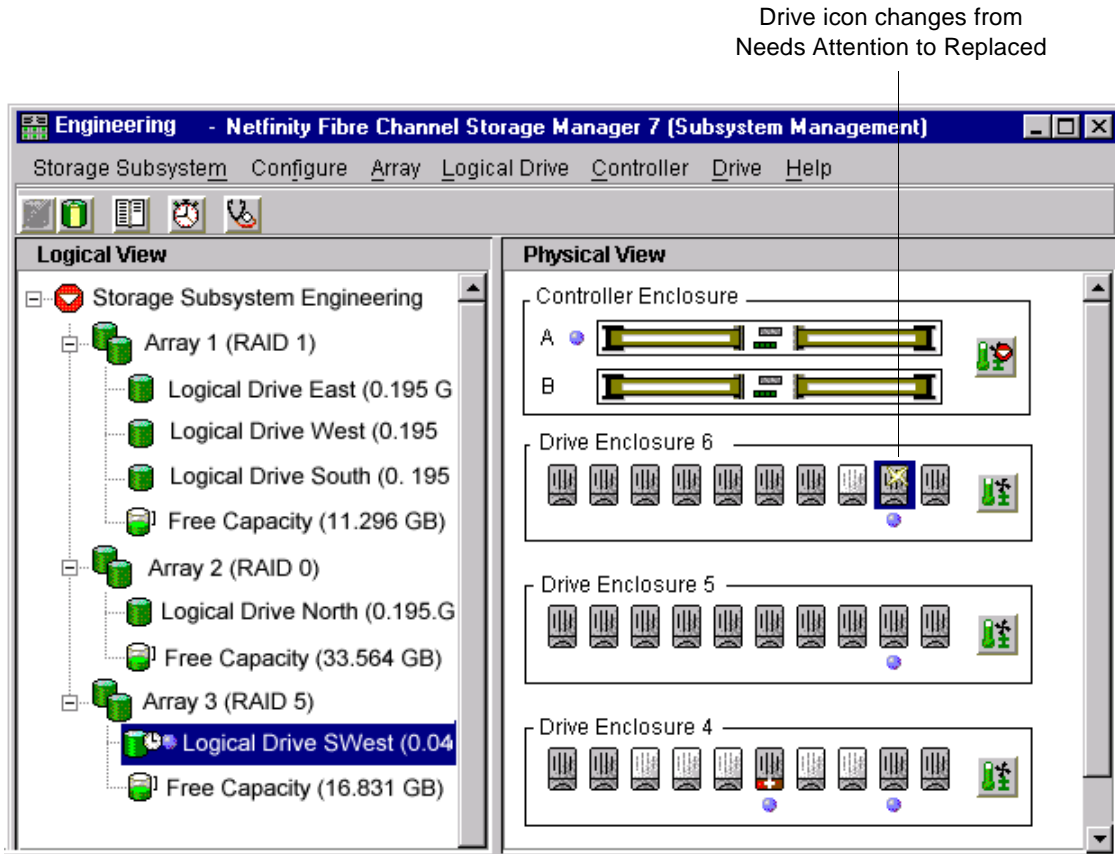


Figure 18. Recovery Guru Window Showing Replaced Status



When the copyback operation is finished, the status icons change to reflect the optimal status of the components, as shown in Figure 19.

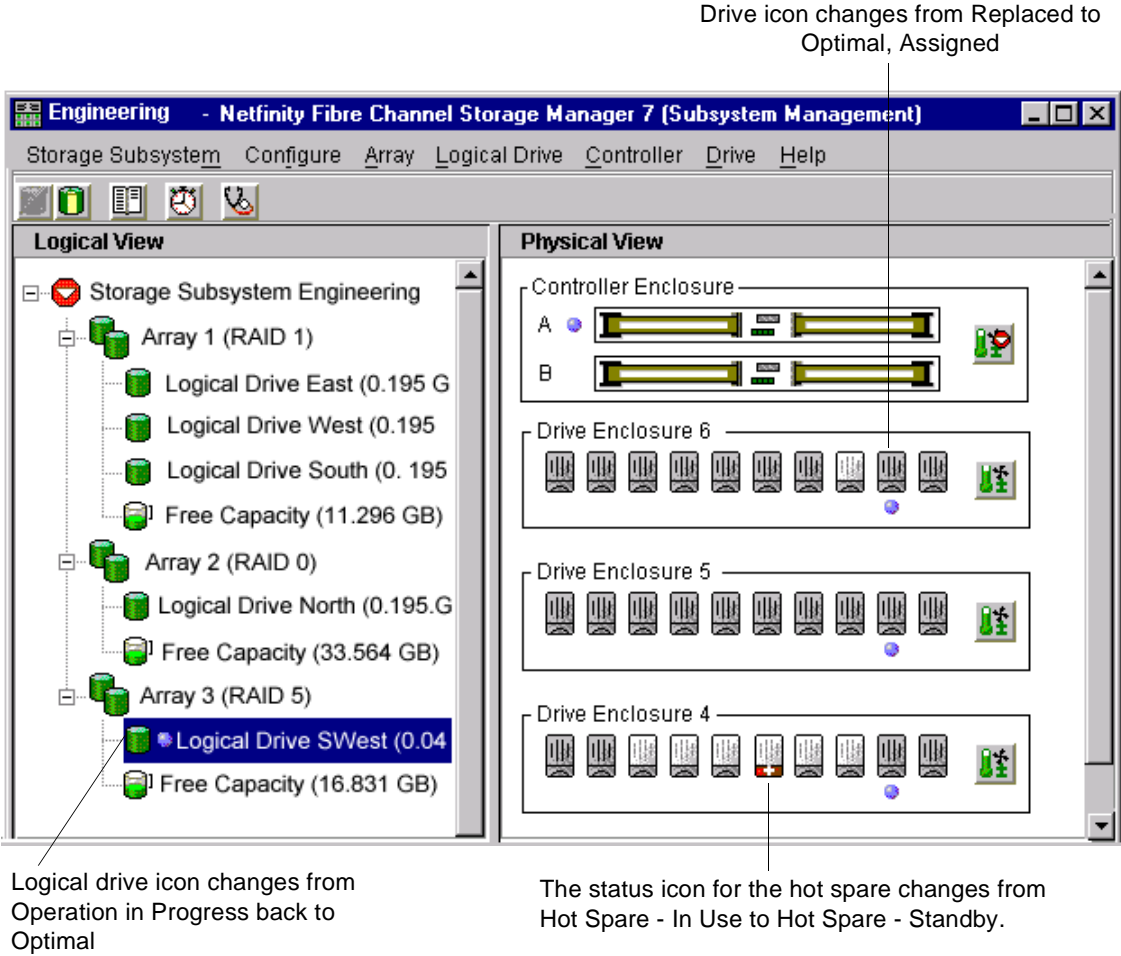


Figure 19. Recovery Guru Window Showing Optimal Status

After you fix the failed battery canister in the controller enclosure, the Components button in the controller enclosure returns to Optimal and the storage subsystem icon returns to Optimal.

# Tuning Storage Subsystems

This section helps the system administrator use data from the Performance Monitor. This section also describes the tuning options available in Netfinity Fibre Channel Storage Manager 7.x, for optimizing storage subsystem and application performance.

Use the Subsystem Management window Performance Monitor to monitor storage subsystem performance in real-time and save performance data to a file for later analysis. You can specify the logical drives and controllers to monitor and the polling interval. Also, you can receive storage subsystem Totals, which is data that combines the statistics for both controllers in an active-active controller pair.

Table 10 describes the Performance Monitor data that displays for selected devices.

**Table 10: Performance Monitor Data Quick Reference**

<b>Data Field</b>	<b>Description</b>
<b>Total I/Os</b>	Total I/Os performed by this device since the beginning of the polling session.  For more information, see "Balancing the I/O Load" on page 33.
<b>Read Percentage</b>	The percentage of Total I/Os that are read operations for this device. Write percentage is calculated as 100 minus this value.  For more information, see "Optimizing the I/O Request Rate" on page 34.
<b>Cache Hit Percentage</b>	The percentage of reads that are processed with data from the cache, rather than requiring a read from disk.  For more information, see "Optimizing the I/O Request Rate" on page 34
<b>Current K/B per second</b>	Average <i>transfer rate</i> during the polling session. The transfer rate is the amount of data in Kilobytes that is moved through the I/O data connection in a second (also called <i>throughput</i> ).  For more information, see "Optimizing the Transfer Rate" on page 33.
<b>Maximum K/B per second</b>	The maximum transfer rate that is achieved during the Performance Monitor polling session.  For more information, see "Optimizing the Transfer Rate" on page 33.
<b>Current I/O per second</b>	The average number of I/O requests serviced per second during the current polling interval (also called an <i>I/O request rate</i> ).  For more information, see "Optimizing the I/O Request Rate" on page 34.

**Table 10: Performance Monitor Data Quick Reference (continued)**

<b>Data Field</b>	<b>Description</b>
<b>Maximum I/O per second</b>	The maximum number of I/O requests serviced during a one-second interval over the entire polling session.  For more information, see “Optimizing the I/O Request Rate.”

## **Balancing the I/O Load**

The Total I/O data field is used for monitoring the I/O activity to a specific controller and a specific logical drive. This field helps you identify possible I/O “hot spots.”

You can identify actual I/O patterns to the individual logical drives and compare those with the expectations based on the application. If a controller has more I/O activity than expected, consider moving an array to the other controller in the storage subsystem using the **Array >> Change Ownership** option.

It is difficult to balance I/O loads across controllers and logical drives because I/O loads are constantly changing. The logical drives and data accessed during the polling session depend on which applications and users are active during that time period. It is important to monitor performance during different time periods and gather data at regular intervals to identify performance trends. The Performance Monitor allows you to save data to a comma-delimited file that you can import to a spreadsheet for further analysis.

If you notice that the workload across the storage subsystem (Total I/O statistic) continues to increase over time while application performance decreases, you might need to add additional storage subsystems to the enterprise.

## **Optimizing the Transfer Rate**

The transfer rates of the controller are determined by the application I/O size and the I/O request rate. In general, a small application I/O request size results in a lower transfer rate, but provides a faster I/O request rate and a shorter response time. With larger application I/O request sizes, higher throughput rates are possible. Understanding the application I/O patterns gives you the maximum I/O transfer rates that are possible for a given storage subsystem.

Because of the dependency on I/O size and transmission media, the only technique you can use to improve transfer rates is to improve the I/O request rate. Use host operating system utilities to gather I/O size data to understand the maximum transfer rates possible. Then use tuning options available in the storage management software to optimize the I/O request rate to reach the maximum possible transfer rate.

## Optimizing the I/O Request Rate

The factors that affect the I/O request rate include:

- I/O access pattern (random or sequential) and I/O size
- Whether write caching is enabled
- Cache hit percentage
- RAID level
- Segment size
- Number of drives in the arrays or storage subsystem
- Fragmentation of files

**Note:** Fragmentation affects logical drives with sequential I/O access patterns, not random I/O access patterns

- Logical drive modification priority

### Determining the I/O Access Pattern and I/O Size

To determine if the I/O has sequential characteristics, try enabling a conservative cache read-ahead multiplier (for example, 4) using the **Logical Drive >> Properties** option. Then examine the logical drive cache hit percentage to see if it has improved. An improvement indicates the I/O has a sequential pattern. For more information, see “Optimizing the Cache Hit Percentage.”

Use host operating system utilities to determine the typical I/O size for a logical drive.

### Enabling Write Caching

Higher write I/O rates are experienced with write caching enabled compared to disabled, especially for sequential I/O access patterns. Regardless of the I/O pattern, be sure to enable write caching to maximize I/O rate and shorten application response time.

### Optimizing the Cache Hit Percentage

A higher Cache Hit Percentage is desirable for optimal application performance and is positively correlated with I/O request rate.

If the Cache Hit Percentage of all logical drives is low or trending downward, and you do not have the maximum amount of controller cache memory installed, this could indicate the need to install more memory.

If an individual logical drive is experiencing a low cache hit percentage, consider enabling cache read-ahead (or prefetch) for that logical drive. Cache read-ahead can increase the Cache Hit Percentage for a sequential I/O workload. If cache read-ahead is enabled, the cache reads the data from the disk. But in addition to the requested data, the cache also fetches more data, usually from adjacent data blocks on the drive. This feature increases the chance that a future request for data is fulfilled from the cache, rather than requiring a disk access.

The cache read-ahead multiplier values specify the multiplier to use for determining how many additional data blocks are read into cache. Choosing a higher cache read-ahead multiplier can increase the cache hit percentage.

If you determine that the I/O has sequential characteristics, try enabling an aggressive cache read-ahead multiplier (for example, 8) using the **Logical Drive >> Properties** option. Then examine the logical drive cache hit percentage to see if it has improved. Continue to customize logical drive cache read-ahead to arrive at the optimal multiplier. (For a random I/O pattern, the optimal multiplier is zero.)

### **Choosing Appropriate RAID Levels**

Use the Read Percentage for a logical drive to determine application behavior. Applications with a high read percentage perform well using RAID level 5 logical drives because of the outstanding read performance of the RAID level 5 configuration.

However, applications with a low read percentage (write-intensive) do not perform as well on RAID 5 logical drives because of the way a controller writes data and redundancy data to the drives in a RAID 5 array. If there is a low percentage of read activity relative to write activity, you can change the RAID level of an array from RAID 5 to RAID 1 for faster performance.

### **Choosing an Optimal Logical Drive Modification Priority Setting**

The modification priority defines how much processing time is allocated for logical drive modification operations versus system performance. The higher the priority, the faster the logical drive modification operations complete, but the slower the system I/O is serviced.

Logical drive modification operations include reconstruction, copyback, initialization, media scan, defragmentation, change of RAID level, and change of segment size.

The modification priority is set for each logical drive using a slider bar on the **Logical Drive >> Properties** dialog box. There are five relative settings on the reconstruction rate slider bar ranging from Low to Highest. The actual speed of each setting is determined by the controller. Choose the Low setting to maximize the I/O request rate. If the controller is idle (not servicing any I/O) it ignores the individual logical drive rate settings and processes logical drive modification operations as fast as possible.

## Choosing an Optimal Segment Size

A *segment* is the amount of data, in kilobytes, that the controller writes on a single drive in a logical drive before writing data on the next drive. Data blocks store 512 bytes of data and are the smallest units of storage. The size of a segment determines how many data blocks it contains. For example, an 8K segment holds 16 data blocks and a 64K segment holds 128 data blocks.

---

**Important:** The segment size was expressed in number of data blocks in previous versions of this storage management software. It is now expressed in Kilobytes.

---

When you create a logical drive, the default segment size is a good choice for the expected logical drive usage. The default segment size is changed using the **Logical Drive >> Change Segment Size** option.

If the I/O size is larger than the segment size, increase the segment size in order to minimize the number of drives needed to satisfy an I/O request. This technique helps even more if you have random I/O access patterns. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If you are using the logical drive in a single-user, large I/O environment such as a multimedia application, storage performance is optimized when a single I/O request is serviced with a single array data stripe (the segment size multiplied by the number of drives in the array used for I/O). In this case, multiple disks are used for the same request, but each disk is only accessed once.

## Defragmenting Files to Minimize Disk Access

Each access of a drive to read or write a file results in movement of the read/write heads. Make sure that the files on the array are *defragmented*. When the files are defragmented, the data blocks making up the files are next to each other, preventing extra read/write head movement when retrieving files. Fragmented files are detrimental to the performance of a logical drive with sequential I/O access patterns.

## Terminology Changes in This Release

Table 11 lists the terminology changes in Netfinity Fibre Channel Storage Manager 7.x. If you are migrating from previous versions of the storage management software or plan to use both versions, use this reference as a guide.

**Table 11: Terminology Change Reference**

<b>Previous Term</b>	<b>New Term</b>
RAID Module	Storage subsystem
LUN*	Logical drive
Drive Group	Array
Drive Module	Drive enclosure
Controller Module	Controller enclosure
ESM CRU	Environmental card canister
Fan CRU	Fan canister
Power Supply CRU	Power supply canister
LED	Indicator light

**Note:** \* Logical Unit Number is still used in this version of the storage management software. It is a logical address that the host uses to access a particular logical drive.





## Chapter 2. Task Reference

---

The Enterprise Management software and Subsystem Management software have unique help systems. This reference is a task-based index to the appropriate help system.

**Table 12: Populating a Management Domain**

<b>See this help system...</b>	<b>For information about this task...</b>
Enterprise Management window	Adding a device to a management domain
	Correcting a partially managed device
	Discovering newly attached host-agent managed storage subsystems
	Performing an initial auto-discovery
	Recovering from corrupted configuration files
	Removing a device from a management domain

**Table 13: Configuring Storage Subsystems**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Assigning a selected unassigned drive as a hot spare drive
	Assigning drives as part of an array
	Consolidating free capacity on an array (defragmenting)
	Creating a logical drive
	Creating a logical drive from free capacity
	Creating a logical drive from unconfigured capacity
	Downloading firmware or NVSRAM
	Expanding the capacity of a selected array by adding unassigned drives
	Increasing the free capacity in a storage subsystem array (deleting a logical drive)

**Table 13: Configuring Storage Subsystems (continued)**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Increasing the unconfigured capacity of a storage subsystem (deleting an array)
	Performing an automatic configuration
	Placing a controller in active or passive mode
	Resetting a storage subsystem configuration
	Returning a selected hot spare drive or drives to an unassigned state
	Specifying logical drive name, usage, desired capacity, controller ownership, and storage partition mapping preference during logical drive creation

**Table 14: Using the Script Editor**

<b>See this help system...</b>	<b>For information about this task...</b>
Enterprise Management window	Adding comments to a script
	Creating logical drives using the Script Editor
	Deleting an array or logical drive using the Script Editor
	Downloading new firmware or NVSRAM to the storage subsystem using the Script Editor
	Editing an existing script
	Executing the currently loaded script
	Interpreting script execution results
	Opening a new script
	Saving the script execution results to a local file
	Saving the script in the Script view
	Using the Script Editor
	Verifying the syntax of the currently loaded script

**Table 15: Configuring Storage Partitions**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Changing a logical drive LUN assignment, host assignment, or host group assignment
	Creating Storage Partitions
	Defining a logical drive-to-LUN mapping
	Deleting a host group, host, or host port from the defined storage subsystem topology
	Deleting a logical drive-to-LUN mapping
	Granting logical drive access to hosts
	Granting logical drive access to host groups
	Moving a host from one host group to another host group
	Moving a host port from one host to another host
	Reconfiguring logical drive-to-LUN mappings
	Renaming a host group, host, or host port
	Replacing a host port after replacing a failed host adapter
	Undefined a host port
	Viewing a list of discovered host ports that are not defined

**Table 16: Protecting Data**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Changing the RAID level of an array
	Checking redundancy information on a selected array
	Configuring a hot spare drive
	Configuring channel protection
	Enabling a media scan on a specific logical drive
	Enabling a redundancy check on an array
	Identifying logical drives that are candidates for a media scan
	Setting the media scan duration
	Specifying when unwritten cache data is written to disk, when a cache flush stops, and the cache block size for a storage subsystem

**Table 17: Event Notification**

<b>See this help system...</b>	<b>For information about this task...</b>
Enterprise Management window	Configuring destination addresses for notifications about an individual storage subsystem
	Configuring destination addresses for notifications about every storage subsystem attached and managed through a particular host
	Configuring destination addresses for notifications about every storage subsystem in the management domain
	Interpreting an e-mail or SNMP trap message
	Specifying management domain global e-mail alert settings
	Validating potential destination addresses
Subsystem Management window	Displaying storage subsystem events in the Event Viewer
	Interpreting event codes
	Interpreting event summary data
	Saving selected events to a file
	Viewing and interpreting event details
	Viewing events stored in the Event Log

**Table 18: Recovering from Problems**

See this help system...	For information about this task...
Subsystem Management window	Failing a selected drive or drives
	Identifying when to use the Recovery Guru
	Initializing drives, logical drives, or arrays
	Interpreting Recovery Guru information
	Manually reconstructing a drive
	Moving arrays (and their associated logical drives) back to their preferred controller owners
	Placing a controller online or offline
	Recovering from connection failures
	Recovering from storage subsystem problems
	Reviving the drives in a selected array or an individual drive
	Saving Recovery Guru information to a text file

**Table 19: Miscellaneous System Administration**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Listing logical or physical components associated with a drive or controller
	Locating a drive, array, or storage subsystem by turning on indicator lights
	Resetting the battery age clock after replacing the battery in the controller enclosure
	Saving storage subsystem information to a text file
	Synchronizing storage subsystem controller clocks with the management station
	Turning off the indicator lights from a Locate operation
	Viewing logical drive data such as logical drive name, worldwide name, status, capacity, RAID level, and segment size
	Viewing a description of all components and properties of a storage subsystem
	Viewing the progress of a logical drive modification operation
	Viewing the properties of a selected drive
	Viewing the properties of a selected controller

**Table 20: Security**

<b>See this help system...</b>	<b>For information about this task...</b>
Subsystem Management window	Changing a storage subsystem password
	Entering a storage subsystem password
Enterprise Management window	Using passwords in the Script Editor

**Table 21: Performance and Tuning**

<b>See this help system...</b>	<b>For information about this task..</b>
Subsystem Management window	Changing segment size on a selected logical drive
	Changing the current and preferred ownership of a selected array
	Changing the polling interval of the Performance Monitor
	Changing the RAID level on a selected array
	Configuring cache block size
	Enabling cache read-ahead
	Interpreting storage subsystem Performance Monitor data
	Modifying the modification priority for a logical drive
	Saving Performance Monitor data to a report
	Saving Performance Monitor data to a spreadsheet
	Selecting logical drives and controllers to monitor with the Performance Monitor
	Specifying cache properties of a logical drive
	Specifying storage subsystem cache settings



# Appendix A. Notices

---

## Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling: (1) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Department 80D, P.O. Box 12195, 3039 Cornwallis, Research Triangle Park, NC 27709, U.S.A. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

## Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM  
Netfinity

Other company, product, and service names may be trademarks or service marks of others.





Part Number: 09N7721

Printed in U.S.A.

December 1999

