

IBM FAStT Storage Manager Version 7.10



# Concepts Guide



IBM FAStT Storage Manager Version 7.10



# Concepts Guide

**Note:** Before using this information and the product it supports, be sure to read the general information under Appendix B, “Notices” on page 69.

**Second Edition (December 2001)**

**© Copyright International Business Machines Corporation 1999, 2001. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Figures</b> . . . . .	v
<b>Tables</b> . . . . .	vii
<b>About this book</b> . . . . .	ix
How this book is organized. . . . .	ix
Notices used in this book . . . . .	x
Related publications . . . . .	x
<b>Chapter 1. Introduction</b> . . . . .	1
Supported machine types . . . . .	1
Terms to know . . . . .	2
New features . . . . .	3
Storage-subsystem components . . . . .	4
Storage-subsystem management. . . . .	5
Direct (out-of-band) management method . . . . .	5
Host-agent (in-band) management method . . . . .	7
Populating a management domain . . . . .	9
Storage-subsystem connectivity. . . . .	10
System requirements. . . . .	11
Hardware requirements . . . . .	11
Reviewing a sample network. . . . .	12
Firmware requirements . . . . .	13
Software components . . . . .	14
Enterprise Management window . . . . .	16
Subsystem Management window Logical and Physical view descriptions . . . . .	17
<b>Chapter 2. Protecting your data</b> . . . . .	19
Using RAID levels for data redundancy . . . . .	19
Protecting data in the controller cache memory . . . . .	20
Configuring hot-spare drives . . . . .	21
Enabling a background media scan . . . . .	21
Fibre Channel I/O data path failover support . . . . .	22
Using available channel protection. . . . .	22
Configuring storage-subsystem password protection . . . . .	22
<b>Chapter 3. Storage-subsystem configuration</b> . . . . .	23
Creating logical drives . . . . .	23
Mappings window overview . . . . .	25
Creating storage partitions. . . . .	26
Storage partitioning . . . . .	28
Storage partitioning example . . . . .	28
Obtaining a feature key. . . . .	29
Creating storage partitions. . . . .	29
Defining storage-partition topology. . . . .	29
Defining logical drive-to-LUN mappings . . . . .	30
Topology-view results . . . . .	31
Mappings-view results . . . . .	32
Heterogeneous hosts overview . . . . .	32
<b>Chapter 4. Maintaining and monitoring storage subsystems</b> . . . . .	35
Maintaining storage subsystems in a management domain. . . . .	35
Storage-subsystem status quick reference . . . . .	36

Failure notification . . . . .	37
Alert notification overview . . . . .	38
Configuring mail server and sender address . . . . .	38
Selecting the node for notification . . . . .	38
Setting alert destinations . . . . .	38
Configuring alert destinations for storage-subsystem critical-event notification . . . . .	38
Event Monitor overview . . . . .	39
Event-monitoring example . . . . .	39
Installing the Event Monitor . . . . .	40
Setting alert notifications . . . . .	40
Synchronizing the Enterprise Management window and Event Monitor . . . . .	41
Recovery Guru window . . . . .	41
Recovering from storage-subsystem failures . . . . .	42
<b>Chapter 5. Tuning storage subsystems . . . . .</b>	<b>47</b>
Balancing the Fibre Channel I/O load . . . . .	47
Optimizing the I/O transfer rate . . . . .	48
Optimizing the Fibre Channel I/O request rate . . . . .	48
Determining the Fibre Channel I/O access pattern and I/O size . . . . .	48
Enabling write-caching . . . . .	49
Optimizing the cache-hit percentage . . . . .	49
Choosing appropriate RAID levels . . . . .	49
Choosing an optimal logical-drive modification priority setting . . . . .	49
Choosing an optimal segment size . . . . .	50
Defragmenting files to minimize disk access . . . . .	50
<b>Chapter 6. Task reference . . . . .</b>	<b>51</b>
<b>Chapter 7. Critical event problem solving . . . . .</b>	<b>57</b>
<b>Appendix A. Getting information, help, and service . . . . .</b>	<b>65</b>
Getting information . . . . .	65
Using the World Wide Web . . . . .	65
Getting information by fax . . . . .	65
Getting help and service . . . . .	65
Using the documentation and diagnostic programs . . . . .	65
Calling for service . . . . .	66
Telephone numbers . . . . .	67
International Warranty Service . . . . .	67
Purchasing additional services . . . . .	68
<b>Appendix B. Notices . . . . .</b>	<b>69</b>
Edition notice . . . . .	69
Trademarks . . . . .	70
Important notes. . . . .	70
<b>Index . . . . .</b>	<b>73</b>

---

## Figures

1. Storage-subsystem physical components . . . . .	4
2. Direct (out-of-band) managed storage subsystems . . . . .	7
3. Host-agent (in-band) managed storage subsystems . . . . .	8
4. Device tree with a management domain . . . . .	9
5. Host-agent (in-band) managed storage subsystem . . . . .	10
6. Direct (out-of-band) managed storage subsystems . . . . .	11
7. Sample network using direct and host-agent managed storage subsystems . . . . .	13
8. Enterprise Management and Subsystem Management windows . . . . .	17
9. Subsystem Management window Logical View and Physical View . . . . .	18
10. Unconfigured and free capacity nodes . . . . .	23
11. Create Logical Drive window . . . . .	24
12. Create Logical Drive - Specify Logical Drive Parameters window . . . . .	24
13. Storage-partitioning example . . . . .	29
14. Topology View . . . . .	30
15. Define New Mapping window . . . . .	31
16. Creating a logical drive-to-LUN mapping in the Topology view . . . . .	31
17. Mappings view . . . . .	32
18. Define New Host Port window . . . . .	33
19. Heterogeneous hosts example . . . . .	33
20. Monitoring storage-subsystem health using the Enterprise Management window . . . . .	35
21. Failure notification in the Subsystem Management window . . . . .	37
22. Event-monitoring example . . . . .	40
23. Location of the Recovery Guru toolbar button . . . . .	42
24. Recovery Guru window . . . . .	43
25. Recovery Guru window showing Replaced status . . . . .	44
26. Recovery Guru window showing Optimal status . . . . .	45





---

## Tables

1. Machine type and supported storage-management software . . . . .	1
2. Old and new terminology . . . . .	2
3. Storage-subsystem physical components . . . . .	4
4. Storage-subsystem logical components . . . . .	4
5. Storage-management architecture hardware components . . . . .	11
6. Storage Manager software components . . . . .	14
7. RAID level configurations . . . . .	19
8. Mappings window views . . . . .	25
9. Storage partitioning terminology . . . . .	26
10. Storage-subsystem status icon quick reference . . . . .	36
11. Performance Monitor tuning options in the Subsystem Management window . . . . .	47
12. Populating a management domain . . . . .	51
13. Configuring storage subsystems . . . . .	51
14. Using the Script Editor . . . . .	52
15. Configuring storage partitions . . . . .	52
16. Protecting data . . . . .	53
17. Event notification . . . . .	53
18. Recovering from problems . . . . .	54
19. <i>Miscellaneous system administration</i> . . . . .	54
20. Security . . . . .	54
21. Performance and tuning . . . . .	55
22. Critical events . . . . .	57



---

## About this book

This book provides the conceptual framework necessary for understanding the IBM® FAStT Storage Manager Version 7.10 for Microsoft® Windows NT®, Windows® 2000, Novell NetWare, and Red Hat Linux® operating-system environments. This *Concepts Guide* is for system administrators. Use this guide to:

- Understand the storage manager software
- Determine what storage-subsystem configuration you will use to maximize data availability
- Set up alert notifications and monitor your storage subsystems in a management domain
- Identify storage manager features that are unique to your specific installation

---

## How this book is organized

Chapter 1, “Introduction” on page 1 provides an introduction of the IBM FAStT Storage Manager Version 7.10 and describes the Enterprise Management window and the Subsystem Management window.

Chapter 2, “Protecting your data” on page 19 discusses software features and how to use the software.

Chapter 3, “Storage-subsystem configuration” on page 23 describes storage-subsystem configuration options that you can use to maximize data availability. This chapter also outlines high-level steps to configure frequently used functions of IBM FAStT Storage Manager.

Chapter 4, “Maintaining and monitoring storage subsystems” on page 35, describes monitoring storage subsystems in a management domain. This chapter also provides the procedure to set up alert notifications that will automatically receive information in the event of failure. The alert notification provides an example of the failure and where to locate the recovery procedure using the Recovery Guru.

Chapter 5, “Tuning storage subsystems” on page 47 discusses tuning options that are available in IBM FAStT Storage Manager Version 7.10.

Chapter 6, “Task reference” on page 51 provides a task-based index to the appropriate online help. There are two separate online help systems in the storage-management software that correspond to each main window: the Enterprise Management window and the Subsystem Management window.

Chapter 7, “Critical event problem solving” on page 57 provides a listing of all critical events that the storage-management software sends in the event of a failure. The event number describes the failure and provides a reference to the procedure that you use to correct the failure.

Appendix A, “Getting information, help, and service” on page 65 provides information about how to obtain help, service, or technical assistance.

Appendix B, “Notices” on page 69 provides product notices and trademark information.

---

## Notices used in this book

This book contains notices to highlight information as follows:

- **Notes**

These notices provide important tips, guidance, or advice.

- **Important**

These notices provide information that might help you avoid inconvenient or problem situations.

- **Attention**

These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage could occur.

---

## Related publications

The following publications are available in Adobe Acrobat Portable Document Format (PDF) on the IBM FAST Storage Manager installation CD and the IBM Web site at <http://www.ibm.com/pc/support/>.

- *IBM FAST Storage Manager Version 7.10 Concepts Guide (this book)*
- *IBM FAST Storage Manager Version 7.10 for Microsoft Windows NT and Windows 2000 Installation and Support Guide*
- *IBM FAST Storage Manager Version 7.10 for Novell NetWare and Red Hat Linux® Installation and Support Guide*
- *IBM FAST Storage Manager Version 7.10 for FAST500 in UNIX® and AIX® Environments Installation and Support Guide*
- *IBM FAST Host Adapter Installation and User's Guide*
- *IBM FAST200 and FAST200 HA Storage Servers Installation and User's Guide<sup>1</sup>*
- *IBM FAST200 Fibre Channel Cabling Instructions<sup>1</sup>*
- *IBM Netfinity FAST500 RAID Controller Enclosure Unit Installation Guide<sup>2</sup>*
- *IBM Netfinity FAST500 RAID Controller Enclosure Unit User's Reference*
- *IBM FAST500 Fibre Channel Cabling Instructions*
- *IBM FAST MSJ User's Guide*
- *IBM Netfinity Fibre Channel RAID Controller Enclosure Unit User's Handbook<sup>3</sup>*

---

1. Printed publication that comes with the IBM FAST200 and FAST200 High Availability (HA) Storage Servers.

2. Printed publication that comes with the IBM FAST500 RAID Controller Enclosure Unit.

3. Printed publication that comes with the IBM Netfinity Fibre Channel RAID controller.

---

## Chapter 1. Introduction

This *Concepts Guide* provides the conceptual framework necessary for understanding the terminology and features of the IBM FAStT Storage Manager Version 7.10 software. Before you begin the installation of the IBM FAStT Storage Manager software, consult the following documentation:

- README.TXT file - Read this text file first. The README file is located in the root directory of the installation CD or refer to the <http://www.ibm.com/pc/support/> IBM Web site for the latest installation and user information about the storage-management software and hardware components.
- IBM FAStT Storage Manager Installation and Support Guide - Use the installation and support guide for your operating system or platform to set up, install, configure, and work with the IBM FAStT Storage Manager Version 7.10 software.

When you finish the entire installation process, refer to the following online help systems, which contain information common to all operating-system environments.

- Enterprise Management window help - Use this online help system to learn more about working with the entire management domain.
- Subsystem Management window help - Use this online help system to learn more about managing individual storage subsystems.

To access the help systems from the Enterprise Management and Subsystem Management windows in IBM FAStT Storage Manager Version 7.1, click **Help** on the toolbar, or press F1.

---

## Supported machine types

Table 1 provides a list of machine type and supported storage-management software.

*Table 1. Machine type and supported storage-management software*

Product name	Machine type	Model	Product release and firmware version	Supported storage-management software version
IBM FAStT200	3542	1RU	4.x	7.02, 7.10
IBM FAStT200 HA	3542	2RU	4.x	7.02, 7.10
IBM Fibre Channel RAID Controller Unit	3526	1RU	3.x, 4.x	6.22, 7.0, 7.01, 7.02, 7.10
IBM FAStT500 RAID Controller Enclosure Unit	3552	1RU	4.x	7.0, 7.01, 7.02, 7.10

**Note:** Product firmware must be at version 4.x to be compatible with Storage Manager Version 7.10.

---

## Terms to know

If you are upgrading from a previous version of the storage-management software, you will find that some of the terms you are familiar with have changed. It is important that you familiarize yourself with the new terminology. Table 2 provides a list of some of the old and new terms.

Table 2. Old and new terminology

Term used in previous versions	New term
RAID module	Storage subsystem
Drive group	Array
Logical unit number (LUN) <sup>1</sup>	LUN
Drive module	Drive enclosure
Controller module	Controller enclosure
Environmental card CRU	Environmental service monitor (ESM) customer replaceable unit (CRU)
Fan canister	Fan CRU
Power-supply canister	Power-supply CRU
LED	Indicator light
Auto-volume transfer	Auto logical-drive transfer
Volume	Logical drive

<sup>1</sup> In Storage Manager 7.10, the term *logical unit number (LUN)* refers to a logical address that is used by the host computer to access a logical drive.

It is important to understand the distinction between the following two terms when reading this document:

### Management station

A management station is a system that is used to manage the storage subsystem. It is attached to the storage subsystem in the following ways.

- Through a TCP/IP Ethernet connection to the controllers in the storage subsystem
- Through a TCP/IP connection to the host-agent software installed on a host computer that is directly attached to the storage subsystem through the Fibre Channel input/output (I/O) path

### Host and host computer

A host computer is a system that is directly attached to the storage subsystem through a Fibre Channel I/O path. This system is used to do the following:

- Serve data (typically in the form of files) from the storage subsystem
- Function as a connection point to the storage subsystem for a remote-management station

### Notes:

1. The terms host and host computer are used interchangeably throughout this book.
2. A host computer can also function as a management station.

---

## New features

The following features are new to the IBM FASiT Storage Manager Version 7.10 software:

- **Event monitoring.** Event monitoring provides the ability to monitor storage subsystems, even when the Enterprise Management window is not open. A new toolbar button and option under **Tools** → **Update Monitor** have been added to the interface to enable updates to the monitor if changes are made to your storage-subsystem configuration.
- **Heterogeneous host computers with storage partitioning.** Heterogeneous host computers with storage partitioning enables you to connect host computers that are running different operating systems to access a single storage subsystem. The host computers can run different operating systems (for example, Solaris and Windows NT) or variants of the same operating system (for example, Windows NT running in a cluster environment, and Windows NT running in a noncluster environment).
- **Configuration replication.** Configuration replication enables you to save your storage subsystem configuration. You can use this feature to replicate a logical configuration from one storage subsystem to another, or to save your storage subsystem configuration for backup purposes.
- **Auto logical-drive transfer.** Firmware version 4.x enables different controllers in the storage subsystem to manage an individual logical drive on a storage subsystem. Previous firmware versions had ownership at the array-level only. If a problem occurs along a path between a host computer and a logical drive, the multipath driver will transfer ownership of the logical drive to the other active controller.
- **Command line.** Use the operating-system command-line interface to issue commands to the storage subsystems. The command-line interface is based on the script processing unit commands found in the Script Editor. Refer to the SM7cli.txt file on the installation CD or on the <http://www.ibm.com/pc/support/> IBM Web site for usage information. The emwdata.bin configuration file is stored in a default directory. The default directory is different, depending on your operating system. Use the operating system Find option to locate the file.
- **Controller diagnostics.** Controller diagnostics are used to test drive-side and host computer Fibre Channel loops from the controller menu in the storage-management software.
- **Access logical drive mapping.** You can assign and change the logical unit number (LUN) assignment for the access logical drive. The access logical drive is used for management communication between the storage subsystem and host computer through the Fibre Channel I/O path.

## Storage-subsystem components

A storage subsystem contains both physical components (such as drives, controllers, fans, and power supplies) and logical components (such as arrays and logical drives). A storage subsystem will span multiple physical enclosures. Figure 1 shows an example of storage-subsystem physical components.




Figure 1. Storage-subsystem physical components

Table 3 describes the storage-subsystems physical components.

Table 3. Storage-subsystem physical components

Components	Description
Drive	An electromagnetic mechanical device that provides the physical data storage media.
Drive enclosure	An enclosure that contains drives, power supplies, fans, environmental service monitors, and other supporting components.
Controller	A system board and firmware that control logical drives and implement the storage-management functions.
Controller enclosure	An enclosure that contains one or more controllers, power supplies, fans, and other supporting components.



Storage-subsystem logical drives provide the physical storage capacity for data. You use the storage-management software to configure the physical capacity into logical components that improve data availability and application performance. Table 4 describes the storage-subsystem logical components.

Table 4. Storage-subsystem logical components

Component	Description
Logical drive	A contiguous subsection of an array that is configured with a RAID level to meet application needs for data availability and I/O performance. A logical drive is seen by the operating system as an unformatted drive.
Storage partition	A logical identity consisting of one or more storage subsystem logical drives that is shared with host computers that are part of a host group or accessed by a single host computer.



Table 4. Storage-subsystem logical components (continued)

Component	Description
Array	A set of physical drives that are grouped together logically by the controllers in a storage subsystem to provide communication from logical drives to host computers.
 Free capacity	A contiguous region of unassigned capacity on a designated array. You will use free capacity to create one or more logical drives.  <b>Note:</b> In the Subsystem Management window Logical view, free capacity is displayed as free-capacity nodes. Multiple free-capacity nodes can exist on an array.
 Unconfigured capacity	The capacity in the storage subsystem from drives that are not assigned to any array. You will use this space to create new arrays.  <b>Note:</b> In the Subsystem Management window Logical view, unconfigured capacity is displayed as an unconfigured-capacity node.

## Storage-subsystem management

IBM FASiT Storage Manager provides two methods for managing storage subsystems: the host-agent (in-band) management method and the direct (out-of-band) management method. Depending on your specific storage-subsystem configurations, you can use either or both methods.

Storage-subsystem management activities include the following:

- Configuring available storage-subsystem capacity into logical drives to maximize data availability and optimize application performance
- Granting access to host computer partitions in the enterprise
- Setting up a management domain
- Monitoring storage subsystems in the management domain for problems or conditions that require attention
- Configuring destinations to receive alert messages for critical events concerning one or more storage subsystems in the management domain
- Recovering from storage-subsystem problems to maximize data availability
- Tuning for optimal application performance

This section describes storage-subsystem Fibre Channel I/O and management connectivity and the hardware and software components in the storage-management architecture. This section also introduces IBM FASiT Storage Manager Version 7.10 and describes how storage subsystems are managed using this software.

### Direct (out-of-band) management method

When you use the direct (out-of-band) management method, you manage storage subsystems directly (out-of-band) over the network through the Ethernet connection to each controller. To manage the storage subsystem through the Ethernet connections, you must define the Internet Protocol (IP) address and host computer name for each controller and attach a cable to the Ethernet connectors on each of the storage subsystem controllers. See Figure 2 on page 7.

Managing storage subsystems using the direct (out-of-band) management method has these advantages:

- The Ethernet connections to the controllers enable a management station running SM7client to manage storage subsystems that are connected to a host computer running an operating system that is supported by IBM FASTT Storage Manager Version 7.10.
- You do not need to use an access volume to communicate with the controllers as you do if you are running the host-agent software. You can configure the maximum number of LUNs that are supported by the operating system and the host adapter that you are using.

Managing storage subsystems using the direct (out-of-band) management method has these disadvantages:

- It requires two Ethernet cables to connect both storage subsystem controllers to a network.
- When adding devices, you must specify an IP address or host computer name for each controller.
- A DHCP/BOOTP server and network preparation tasks are required. For a summary of the preparation tasks, refer to the *installation and support guide* for your operating system.

Figure 2 shows a system in which storage subsystems are using the direct (out-of-band) management method.

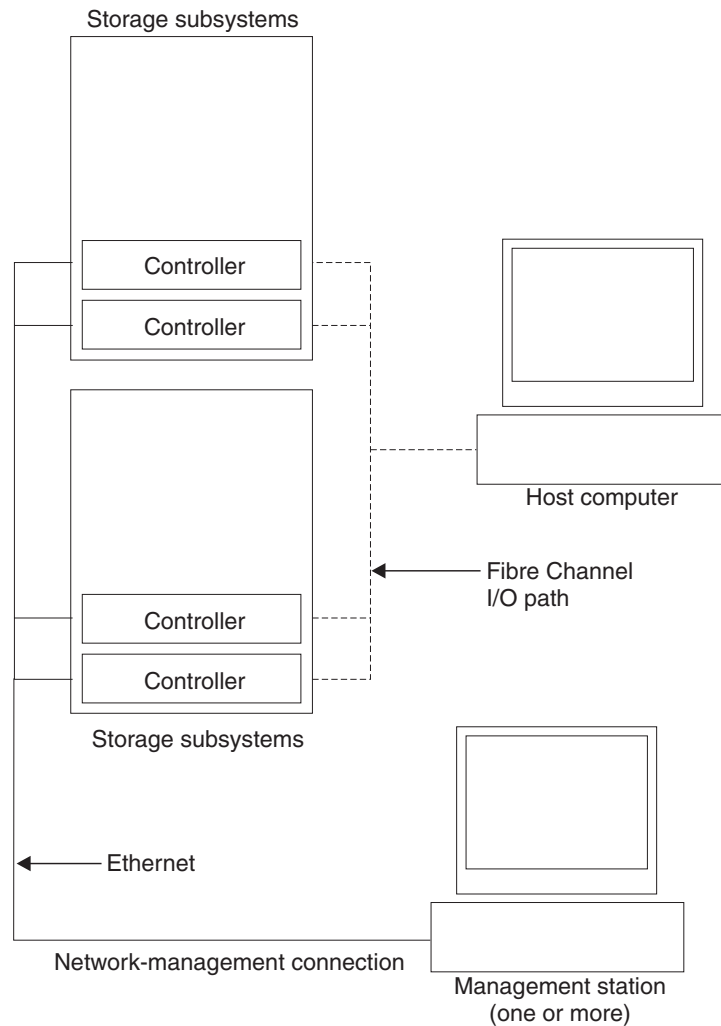


Figure 2. Direct (out-of-band) managed storage subsystems

## Host-agent (in-band) management method

When you use the host-agent management method, the controllers in the storage subsystem are managed through the host-agent (in-band) Fibre Channel network connection to a host computer, rather than through the direct (out-of-band) Ethernet network connections to each controller. The host-agent software on the host computer enables communication between the management software and the controllers in the storage subsystem. To manage a storage subsystem using the host-agent management method, you must install the host-agent software on the host computer and then use the Enterprise Management window to add the host computer to the management domain. By including the host computer in the domain, you will include attached host-agent managed storage subsystems also.

Managing storage subsystems through the host agent has these advantages:

- You do not have to run Ethernet cables to the controllers.
- You do not need a DHCP/BOOTP server to connect the storage subsystems to the network.

- You do not need to perform the controller network configuration tasks.
- When adding devices, you must specify a host computer name or IP address *only* for the host computer instead of for the individual controllers in a storage subsystem. Storage subsystems that are attached to the host computer are automatically detected.

Managing storage subsystems through the host agent has these disadvantages:

- You are limited to configuring one fewer LUN than the maximum number that is allowed by the operating system and the host adapter that you are using.
- The host agent requires a special logical drive, called an *access volume*, to communicate with the controllers in the storage subsystem.
- If you are upgrading controllers from firmware version 3.x to version 4.x and your host computer has already configured its maximum number of LUNs, you must give up a LUN to be used as an access volume.

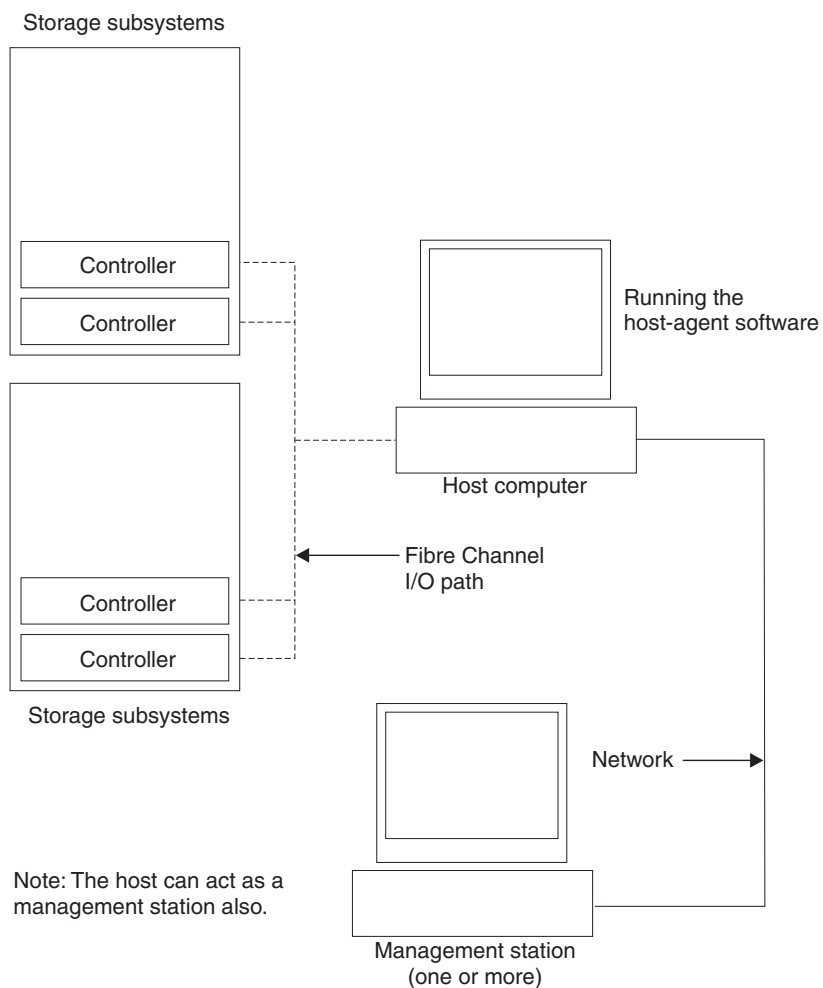


Figure 3. Host-agent (in-band) managed storage subsystems

---

## Populating a management domain

A *management domain* is a collection of storage subsystems that you want to manage. The management domain is displayed in the Enterprise Management window device tree as shown in Figure 4.




Figure 4. Device tree with a management domain

There are two ways to populate a management domain:

- **Using the Automatic Discovery option.** From the Enterprise Management window, select the Automatic Discovery option to automatically discover direct managed and host-agent managed storage subsystems on the local subnetwork and add them to the management domain. The storage-management software discovers host-agent managed storage subsystems by first discovering the host computers that provide host-agent network-management connections to the storage subsystems. Then, the host computer and associated storage subsystems appear in the device tree.
- **Using the Add Device option.** From the Enterprise Management window, select the Add Device option if you want to directly manage the storage subsystem. Type a host computer name or IP address for each controller in the storage subsystem. For a host-agent managed storage subsystem, type a name or IP address for the host computer that is attached to the storage subsystem.

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

## Storage-subsystem connectivity

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

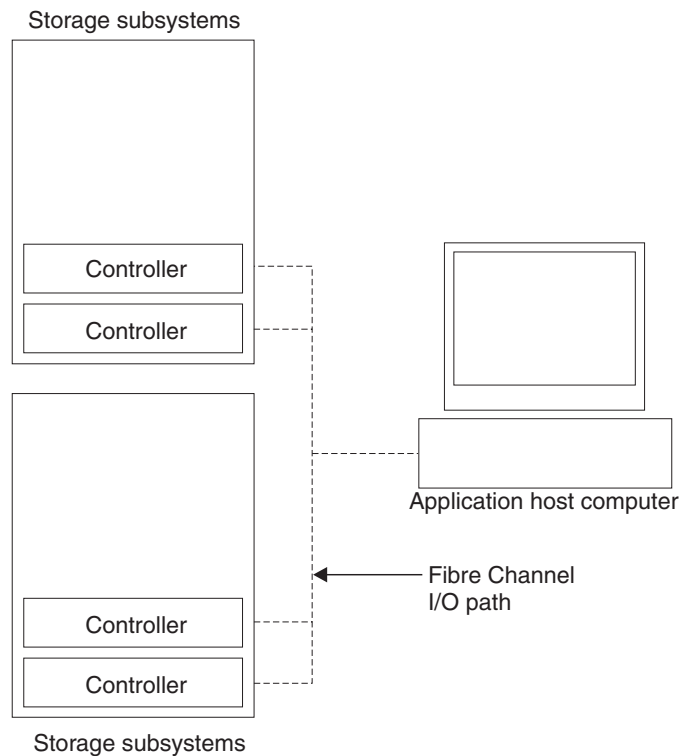


Figure 5. Host-agent (in-band) managed storage subsystem

When using previous versions of the storage-management software, the software is installed on the host computer. Storage-management commands are sent from the host computer to the storage-subsystem controllers through the Fibre Channel I/O path (referred to as *in-band* management).

In addition to the host computer, the IBM FAStT Storage Manager Version 7.10 is installed on management stations from anywhere on the corporate network. Storage-management commands are sent from the management station through an Ethernet network to the storage subsystem (referred to as *out-of-band* management). The connection between the management station and the storage-subsystem controller is called the network-management connection. The two kinds of network management connections are described in “Storage-subsystem management” on page 5.

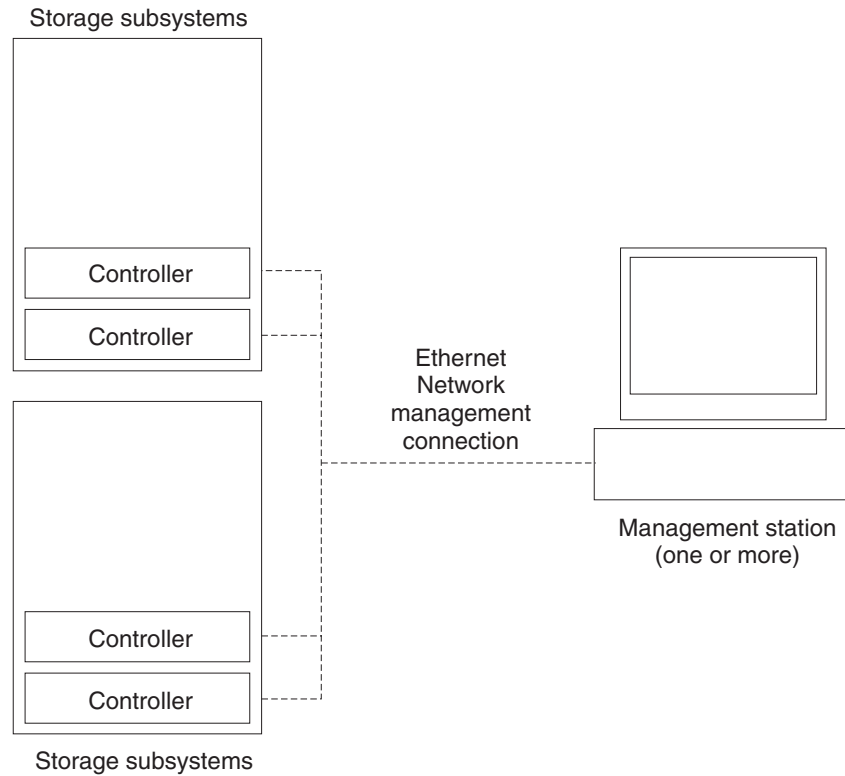


Figure 6. Direct (out-of-band) managed storage subsystems

## System requirements

This section provides detailed information about the hardware, software, and storage-management architecture for IBM FASTT Storage Manager Version 7.10.

## Hardware requirements

Table 5 lists the hardware required for installing Storage Manager 7.10.

Table 5. Storage-management architecture hardware components

Component	Description
Management station (one or more)	A computer that is connected through an Ethernet cable to the host computer or directly to the storage controller.
Network-management station (optional, for SNMP traps)	A computer with installed SNMP-compliant network-management software that receives and processes information about managed network devices using Simple Network Management Protocol (SNMP).  The storage-management software sends critical alerts (using SNMP trap messages) to configured destinations.
DHCP/BOOTP or BOOTP-compatible server (for only direct-managed storage subsystems)	A server used to assign the network-specific information such as an internet protocol (IP) address and host computer name for each controller.  <b>Note:</b> This server is not required if you are managing all storage subsystems through the Fibre Channel I/O path using a host agent.

Table 5. Storage-management architecture hardware components (continued)

Component	Description
Host computer	A computer running one or more applications that accesses the storage subsystem through the Fibre Channel I/O data connection.
Storage subsystem (one or more)	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).
File server	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.

## Reviewing a sample network

Figure 7 on page 13 shows an example of a direct (out-of-band) managed storage subsystem network setup. Network A contains the following components:

- DHCP/BOOTP server
- Network Management Station (NMS) for Simple Network Management Protocol (SNMP) traps
- Host computer that is connected to a storage subsystem through a Fibre Channel I/O path
- Management station connected by Ethernet cable to the storage subsystem controllers

Figure 7 on page 13 shows an example of a host-agent managed storage subsystem network setup. Network B contains the following components:

- A host computer that is connected to a storage subsystem through a Fibre Channel I/O path
- A management station that is connected by Ethernet cable to the host computer



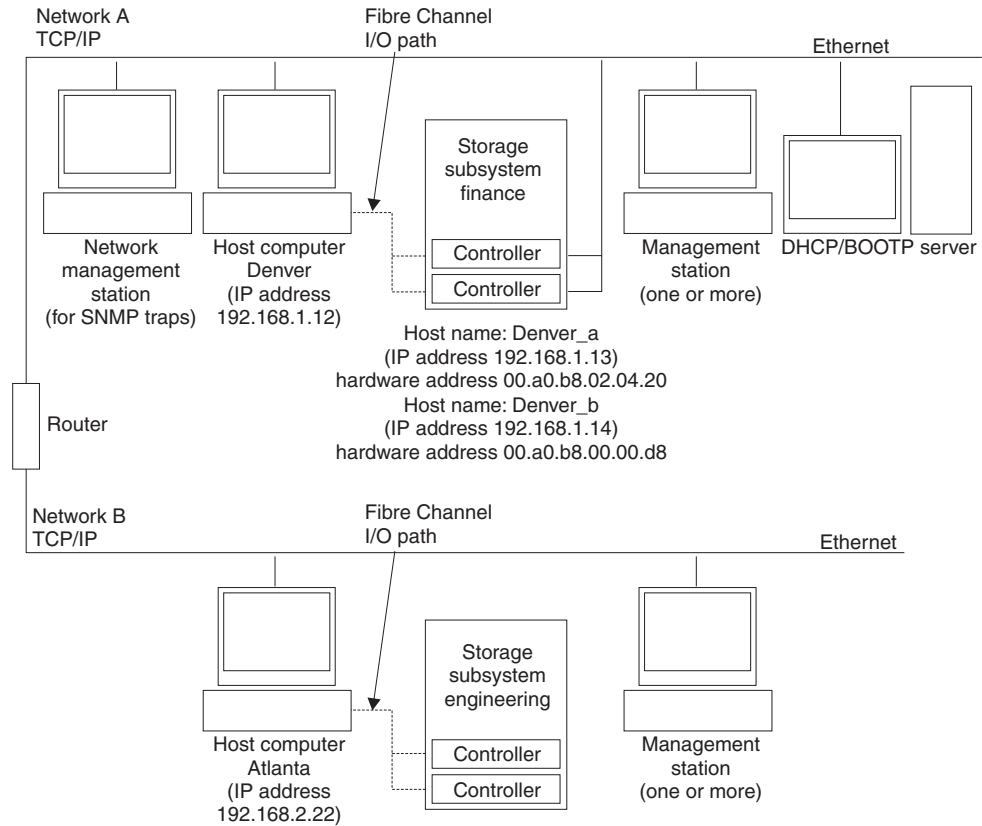


Figure 7. Sample network using direct and host-agent managed storage subsystems

## Firmware requirements

Storage Manager 7.10 operates only with controller machine types 3526, 3552, or 3542, and firmware version 4.00 through 4.01. If you want to manage controllers with Storage Manager 7.10, you must upgrade the firmware to version 4.00 through 4.01. Refer to the appropriate IBM FAStT Storage Manager Installation and Support Guide for more information.

---

## Software components

Table 6 describes the IBM FASiT Storage Manager Version 7.10 software components.

Table 6. Storage Manager software components

Component	Description
Storage Manager 7.10 client (SM7client)	<p>The Storage Manager 7.10 client (SM7client) component provides the graphical user interface (GUI) for managing storage subsystems through the Ethernet network or from the host computer. The SM7client contains two main components:</p> <ul style="list-style-type: none"><li>• <b>Enterprise Management.</b> You can use the Enterprise Management component for adding, removing, and monitoring storage subsystems within the management domain.</li><li>• <b>Subsystem Management.</b> You can use the Subsystem Management component for managing the components of an individual storage subsystem.</li></ul> <p>The Storage Manager 7 client is called <i>thin</i> because it only provides an interface for storage management based on information supplied by the storage-subsystem controllers. When you install the SM7client software component on 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 the SM7client.</p>

Table 6. Storage Manager software components (continued)

Component	Description
Storage Manager 7.10 agent Host agent (SM7agent)	<p>The Storage Manager 7.10 agent (SM7agent) package contains the host-agent software. You can use the host-agent software to manage storage subsystems through the host computer Fibre Channel I/O path. The host-agent software receives requests from a management station that is connected to the host computer through a network connection and passes the requests to the storage-subsystem controllers through the Fibre Channel I/O path.</p> <p>The host agent, along with the network connection on the host computer, provides an in-band host agent type network-management connection to the storage subsystem instead of the out-of-band direct network management connection through the individual Ethernet connections on each controller.</p> <p>The management station can communicate with a storage subsystem through the host computer that has host-agent management software installed. The host agent receives requests from the management station through the network connection to the host computer and sends them to the controllers in the storage subsystem through the Fibre Channel I/O path.</p> <p><b>Note:</b> Host computers 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.</p> <p>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 computer with the host-agent software installed. In this case, you can remove the duplicate storage subsystem icon from the device tree by using the Remove Device option in the Enterprise Management window.</p>
Redundant Disk Array Controller (RDAC) multipath driver software	<p>RDAC is a Fibre Channel I/O path failover driver that is installed on host computers. Usually, a pair of active controllers is located in a storage subsystem enclosure. Each logical drive in the storage subsystem is assigned to a controller. The controller is connected to the Fibre Channel I/O path between the logical drive and the host computer through the Fibre Channel network. When a component in the Fibre Channel I/O path fails, such as a cable or the controller itself, the RDAC multipath 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 Fibre Channel I/O path failover drivers and do not require this multipath driver.</p>

Table 6. Storage Manager software components (continued)

Component	Description
The IBM Storage Area Network (IBMSAN) driver	The IBM Storage Area Network (IBMSAN) driver is used in a Novell NetWare environment to provide multipath I/O support to the storage controller. The IBMSAN driver is a custom device module (CDM) that is installed with the IBM FASTT Host Adapter driver. In case of a failure along the I/O path, the driver will send the I/O requests through the alternate path and auto logical-drive transfer will move the logical drive within the FASTT Storage Server or controller enclosure.
Storage Manager 7 Utility (SM7util)	Use the Storage Manager 7 Utility software package to register and map new logical drives to the operating system. The Storage Manager 7 Utility is installed on all host computers. The host computers are attached to the storage subsystem through the Fibre Channel connection. The Storage Manager 7 Utility software package contains the following two components: <ul style="list-style-type: none"> <li>• <b>Hot Add utility.</b> The Hot Add utility enables you to register newly created logical drives with the operating system.</li> <li>• <b>SM7devices utility.</b> You can use the SM7devices utility to associate storage subsystem logical drives with operating-system device names.</li> </ul> <p><b>Note:</b> Refer to the storage-management software README file for all supported operating systems.</p>
Microsoft Virtual Machine (MSVM)	Microsoft Virtual Machine is required to support the Event Monitor option of the SM7client package on a host computer running Windows NT. You must manually install Microsoft Virtual Machine on your management station, host computer, or server as applicable to your configuration.

## Enterprise Management window

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

- Add and discover the storage subsystems that 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

The SM7client Graphical User Interface (GUI) displays two main windows: the Enterprise Management window and the Subsystem Management window, as shown in Figure 8 on page 17.

The emwdata.bin configuration file contains a list of the storage subsystems that are included in the management domain. After adding the storage subsystems, use the Enterprise Management window primarily for course-level monitoring and alert notification of non-optimal storage-subsystem conditions. Also, you can use it to open the Subsystem Management window for a particular storage subsystem. The emwdata.bin configuration file is stored in a default directory. The default directory is different, depending on your operating system.

**Note:** If multiple users have installed the SM7client in a Windows NT environment, there will be multiple emwdata.bin files throughout the system.

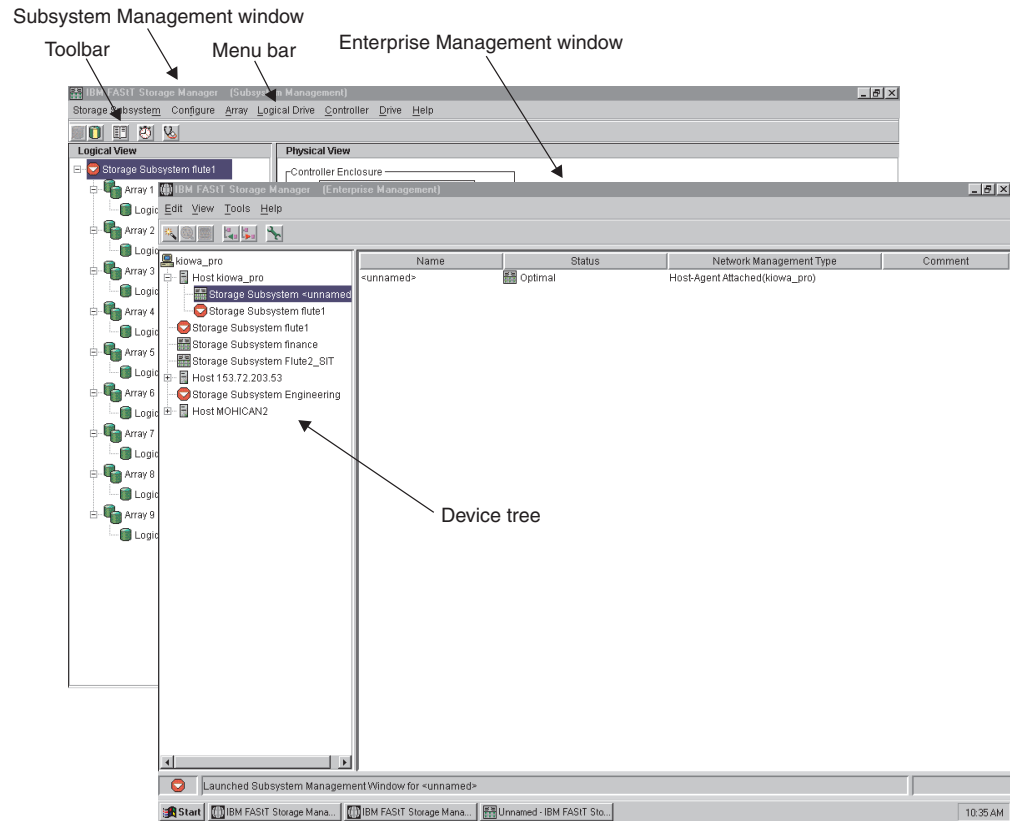


Figure 8. Enterprise Management and Subsystem Management windows

## Subsystem Management window Logical and Physical view descriptions

The Logical view is displayed in the left pane of the Subsystem Management window. The storage-subsystem capacity is organized into arrays and logical drives. The Physical view is displayed in the right pane of the Subsystem Management window. The Physical view displays the physical devices in the storage subsystem, such as controllers and drives. When you click a logical drive or other item in the Logical view, the associated physical components are displayed in the Physical view, as shown in Figure 9 on page 18.

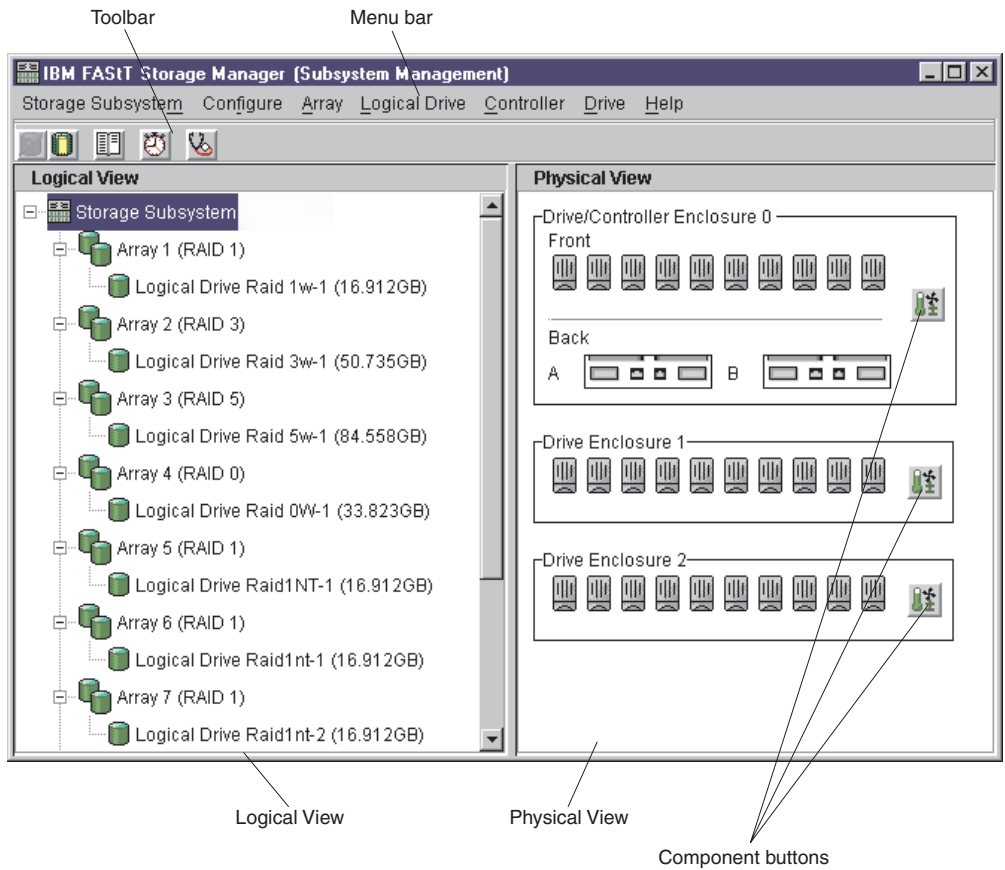


Figure 9. Subsystem Management window Logical View and Physical View

---

## Chapter 2. Protecting your data

When configuring a storage subsystem, review the appropriate data protection strategies and decide how you will organize the storage capacity into logical drives that are shared among hosts in the enterprise.

Storage subsystems are designed for reliability, maximum data protection, and twenty-four hour data availability through a combination of hardware redundancy and controller firmware configurations.

The examples of hardware redundancy are:

- Dual hot-swap RAID controller units
- Dual hot-swap fans
- Dual hot-swap power supplies
- Internal battery unit to protect cache memory in the event of power outages
- Dual Fibre Channel drive loops from the controller enclosure to all of the Fibre Channel enclosures
- Reserve hot-spare drives

The examples of software configurations are:

- Support for different logical-drive RAID levels
- Orthogonal RAID scripting support
- Multiple write-caching options and the ability to set thresholds for the cache-flushing algorithm
- Hot-spare drive swapping configuration
- Background media scan
- Storage subsystem managed password protection
- Fibre Channel I/O path to the controller failover

---

### Using RAID levels for data redundancy

Redundant Array of Independent Disks (RAID) relies on a series of configurations, called levels, to determine how user and redundancy data is written and retrieved from the drives. The IBM FAStT Storage Manager software offers four RAID level configurations: RAID level-0, level-1, level-3, and level-5. Each level provides different performance and protection features.

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

Table 7 on page 20 describes the RAID level configurations available with IBM FAStT Storage Manager Version 7.10.

Table 7. RAID level configurations

RAID level	Description
RAID 0	RAID level-0 offers simplicity, but does not provide data redundancy. A RAID level-0 array spreads data across all drives in the array. If one drive in the array fails, all logical drives contained in the array fail.
RAID 1	<p>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 level-1 array, the mirrored drive takes over.</p> <p>RAID level-1 is costly in terms of capacity. One-half of the drives are used for redundant data.</p>
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 amounts of sequential data. In these applications, the I/O size is large, and all 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 file-system 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. The other drives are available to concurrently service other I/O read requests and deliver fast read I/O request rates.</p>

**Note:** One array uses a single RAID level 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 enables the controller cache memory to store write operations from the host computer, improving system performance. However, a controller can fail with user data in its cache that has not been transferred to the logical drive. 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 enables 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. That is, if one controller fails, the other will complete all outstanding write operations. You can enable the write-cache mirroring parameter for each logical drive.

To prevent data loss or damage, the controller writes cache data to the logical drive periodically. When the cache holds a specified start percentage of unwritten data, the controller writes the cache data to the logical drive. When the cache is flushed



down to a specified stop percentage, the flush is stopped. For example, you can specify that the controller start flushing the cache data when the cache reaches 80% full and stop writing 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 computer 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 logical drive 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 battery CRUs at the recommended time intervals (refer to your Storage Manager Installation Guide for product specific replacement instructions). The storage-management software features a battery-age clock that 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.

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 removed.

**Attention:** For maximum data availability, do not enable the write-caching without batteries 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 assign available physical 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 level-1, level-3, or level-5 array. If the logical drive in a RAID level-1, level-3, or level-5 array fails, the controllers automatically use a hot-spare drive to replace the failed logical drive while the storage subsystem is operating. The controller uses redundancy data to automatically reconstruct the data from the failed logical drive to the replacement (hot-spare) drive. This is called *reconstruction*.

The hot-spare drive adds another level of redundancy to the storage subsystem. If a logical drive fails in the storage subsystem, the hot-spare drive is automatically substituted without requiring a physical swap. If the hot-spare drive is available when a logical drive fails, the controller uses redundancy data to reconstruct the data from the failed logical drive to the hot-spare drive. When you have physically replaced the failed logical drive, the data from the hot-spare drive 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 logical-drive read and write functions. The media-scan process scans all logical-drive data to verify that it is accessible, and optionally scans the logical drive redundancy data.

---

## Fibre Channel I/O data path failover support

If the operating system on the application host computer does not include a multipath failover driver, install the RDAC multipath driver that comes with the storage-management software.

The RDAC multipath driver manages the Fibre Channel I/O path failover process for storage subsystems with redundant controllers. If a component (for example, a cable, controller, or host adapter) fails along the Fibre Channel I/O path, the RDAC multipath driver automatically reroutes all I/O operations to the other controller.

**Note:** The RDAC multipath driver is not supported on all operating systems. Refer to your *Storage Manager Installation Guide* for more information.

The IBM Storage Area Network (SAN) driver is used in a Novell NetWare environment to provide multipath I/O support to a storage server or controller enclosure. The IBMSAN driver is a custom driver module (CDM) that is installed with the IBM FASiT Host Adapter driver. In case of a failure along the I/O path, the driver will send the I/O requests through the alternate path and Auto-Volume Transfer (AVT) will move the logical drive within the FASiT Storage Server. Refer to the <http://www.ibm.com/pc/support> IBM Web page for the current README file, SAN driver, and installation instructions.

---

## Using available channel protection

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

For added security, you can configure a password for each storage subsystem you manage by clicking **Storage Subsystem** → **Change Password**.

After you have set the password for each storage subsystem, you will be prompted for that password the first time you attempt a destructive operation in the Subsystem Management window. You are asked for the password only once during a single management session.

## Chapter 3. Storage-subsystem configuration

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

### Creating logical drives

This section provides a basis for understanding the creation of logical drives. For detailed help, refer to the Subsystem Management window online help.

A logical drive is a logical structure you create on a storage subsystem for data storage. A logical drive is defined by a set of physical drives called an array, which has a defined RAID level and capacity. You can define logical drives from unconfigured-capacity or free-capacity nodes on the storage subsystem.

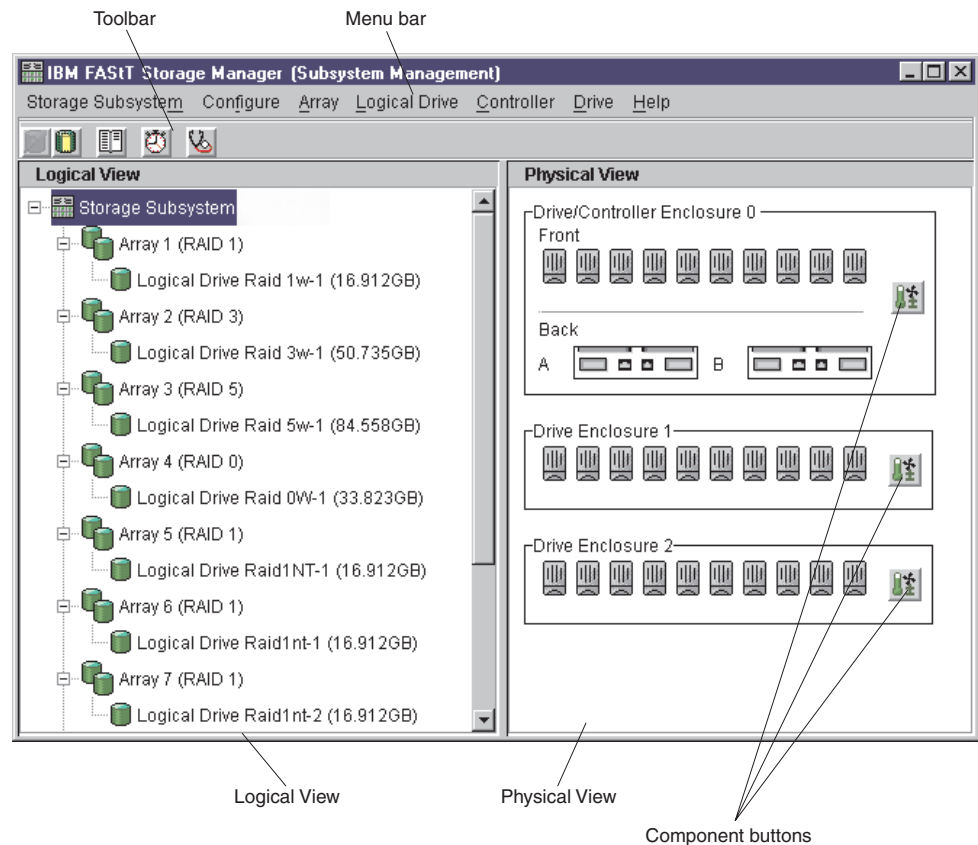


Figure 10. Unconfigured and free capacity nodes

You can configure logical drives from either unconfigured capacity or free capacity in a storage subsystem from the Subsystem Management window. If you have not configured any logical drives on the storage subsystem, the only node available is the unconfigured-capacity node.

Perform the following steps to create logical drives from unconfigured or free capacity from the Subsystem Management window.

1. To configure logical drives from the unconfigured-capacity node, click **Unconfigured-Capacity Node**, and then, click **Configure** → **Create Array** → **Logical Drive**.

The Create Logical Drive window opens.

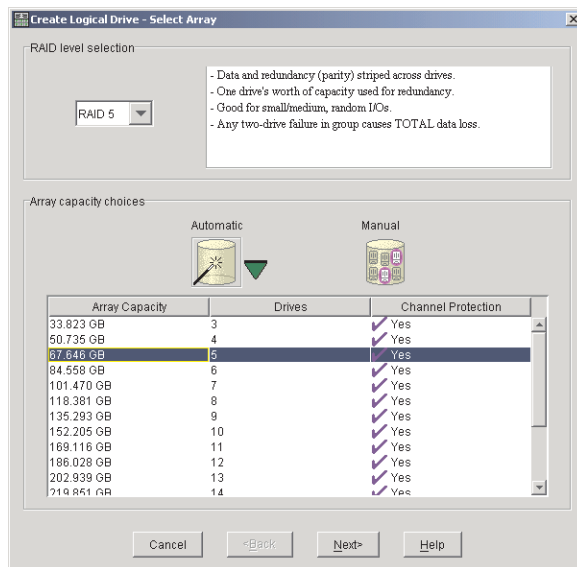


Figure 11. Create Logical Drive window

2. From the Create Logical Drive window, specify the following two array parameters:
  - RAID level
  - Array capacity (how many drives you want to include in the array)
3. Click **Next**.

The Create Logical Drive - Specify Logical Drive Parameters window opens.

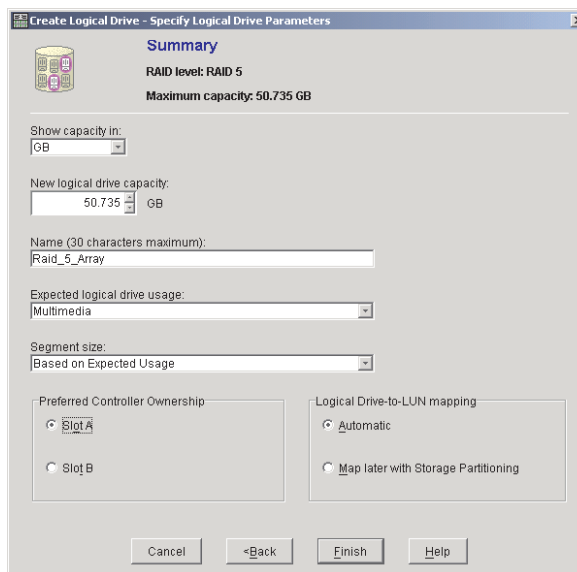


Figure 12. Create Logical Drive - Specify Logical Drive Parameters window

4. Specify the logical-drive parameters for the first logical drive on the array. Parameters include capacity, logical drive name, expected usage, segment size, controller ownership, and logical drive-to-LUN mapping preference.  
The capacity parameter is the capacity that you want for the individual logical drive, not the array. Choose to use the entire capacity that is remaining, or specify a smaller capacity.  
The expected-usage parameter is a convenient way of selecting an optimal segment size for the application that will use this logical drive. From the Subsystem Management window, click **File** → **System** → **Database** → **Multimedia** and then, select **Based on Expected Usage for Segment Size**, to automatically assign the segment size that is optimal for the expected usage.
5. Set the logical drive-to-LUN mapping preference to specify whether you will be using storage partitioning. Use one of the following two settings:
  - **Automatic.** This setting specifies that a LUN is automatically assigned to the logical drive using the next available LUN within the default host group. This setting grants logical-drive access to host groups or host computers that have no specific logical drive-to-LUN mappings (designated by the default host group node in the Topology view). If you do not use storage partitioning, you must specify the Automatic setting.
  - or
  - **Map later with Storage Partitioning.** This setting specifies that a LUN not be assigned to the logical drive during creation. This setting enables definition of a specific logical drive-to-LUN mapping and creation of storage partitions using the **Configure** → **Storage Partitioning** option. When you use storage partitioning, specify this setting.

**Note:** To add free capacity to an array using the Subsystem Management window, click **Array** → **Add Free Capacity**.
6. Click **Finish** to create the array and logical drive.
7. Register the logical drive with the operating system.  
After you create logical drives with automatic logical drive-to-LUN mappings, follow the appropriate instructions in the Installation and Support Guide to enable the operating system to discover the new logical drive.
8. If you selected the free-capacity node, you will create logical drives in the existing array. Click **Configure** → **Create Array** → **Logical Drive**.  
The Create Logical Drive - Specify Logical Drive Parameter window opens (see Figure 12 on page 24).

## Mappings window overview

Before you use storage partitions, you must define topological elements. To define topological elements, use the **Configure** → **Storage Partitioning** option in the Subsystem Management window to display the Mappings window. Use the Mappings window to define the storage-topology elements (for example, host groups, host computers, or host ports), to define logical drive-to-LUN mappings, and to view storage partitioning and heterogeneous host computer information. The Mappings window has two views, Topology and Mappings, shown in Figure 14 on page 30 and Figure 17 on page 32. Table 8 on page 26 describes the two Mappings window views.

Table 8. Mappings window views

View	Description
Topology	<p>The Topology view shows defined topological elements (host ports, host computers, 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 host computers. You will use the Mappings view to show the default LUNs used to access logical drives in the storage subsystems.</p>
Mappings	<p>The Mappings view shows the default and specific logical drive-to-LUN mappings in a storage subsystem. In a table format, information displayed about the logical drives includes the following:</p> <ul style="list-style-type: none"> <li>• The topological entity that can access the logical drive</li> <li>• A logical-drive name</li> <li>• A unique-logical drive identifier</li> <li>• The LUN used to access the logical drive</li> </ul>

## Creating storage partitions

You can use the Storage Partitions feature of the IBM FASTT Storage Manager software to consolidate logical drives into sets called *storage partitions*. You grant visibility of partitions to defined host computers or a defined set of hosts called a *host group*. Storage partitions enable host computers to share storage capacity. Storage partitions consolidate storage and reduce storage-management costs. Table 9 describes the storage partitioning terminology that is used in the Mappings window.

Table 9. Storage partitioning terminology

Term	Description
Storage partition	Storage subsystem logical drives that are visible to a host computer or are shared among host computers that are part of a host group.
Storage-partition topology	The Topology view of the Mappings window displays the default host group, the defined host group, host computer, and host-port nodes. You must define the host port, host computer, and host group topological elements to grant access to host computers and host groups using logical drive-to-LUN mappings.
Host port	<p>Host ports physically reside on the host adapters and are automatically discovered by the IBM FASTT Storage-Manager software. To give a host computer access to a partition, you must define its associated host ports.</p> <p>The host ports request data from a logical drive on behalf of the host computer; therefore, without associated host ports, a host computer 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 9. Storage partitioning terminology (continued)

Term	Description
Host computer	<p>An entity in the storage-partition topology that is a defined, logical set of host ports that need to access a logical drive.</p> <p><b>Note:</b> A defined host computer corresponds to a single computer running one or more applications that accesses a storage subsystem. A host computer must not belong to a defined host group unless the host computer must share access to a partition with other host computers.</p>
Host group	<p>An entity in the storage-partition topology that defines a logical collection of host computers that require shared access to one or more logical drives.</p> <p><b>Note:</b> You can define a host group that corresponds to a cluster or set of host computers that provide failover support. Host computers in a defined host group are 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 computer in a host group.</p>
Default host group	<p>A logical collection of discovered host ports, defined host computers, and defined host groups in the storage-partition 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>
LUN	<p>Logical unit number. The number a host computer uses to access a logical drive. Each host computer has its own LUN address space.</p>
Specific logical drive-to-LUN mapping	<p>An association of a logical drive with a single LUN. When you create a specific logical drive-to-LUN mapping, you specify both the LUN that is used to access the logical drive and the defined host computer or host group that can access the logical drive.</p>
Default logical drive-to-LUN mapping	<p>This mapping enables host groups or host computers that do not have specific logical drive-to-LUN mappings (host computers 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 legacy logical drives (created using previous versions of the IBM FAStT Storage Manager software) are automatically given default logical drive-to-LUN mappings.</p> <p>After installing IBM FAStT Storage Manager Version 7.10 software, specific logical drive-to-LUN mappings are created.</p>



Table 9. Storage partitioning terminology (continued)

Term	Description
Storage-partitions mapping preference	<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>

## Storage partitioning

You can use storage partitioning to enable access to logical drives by designated host computers in a host group or by a single host computer. A storage partition is created when a collection of host computers (a host group) or a single host computer is associated with a logical drive-to-LUN mapping. The mapping defines which host group or host computer can access a particular logical drive in a storage subsystem. Host computers and host groups can access data only through assigned logical drive-to-LUN mappings. See “Defining logical drive-to-LUN mappings” on page 30 for more information. For detailed information about storage partitions, see the Subsystem Management window online help.

## Storage partitioning example

In the example shown in Figure 13 on page 29, four host computers (Omaha A, Omaha B, KC-A and KC-B) are connected to storage subsystem Midwest. Host computer KC-B has exclusive access to logical drive Financial. Host computers KC-A and KC-B have access to logical drives Legal and Engineering, and Omaha A and Omaha B have access to logical drives Marketing and HResources.

In Figure 13 on page 29, three storage partitions are configured. The first storage partition consists of logical drive Financial. This logical drive is accessed by host computer KC-B using LUN 5. Even though host computer KC-B is part of the logical host group Kansas City, host computer KC-A cannot access this logical drive because the logical drive-to-LUN mapping is created with host computer KC-B, rather than the host group Kansas City.

The second storage partition consists of logical drives Legal and Engineering. This logical drive-to-LUN mapping is created with the host group Kansas City. For example, the logical drive is accessed by host computer KC-A in host group Kansas City using LUNs 2 and 4, respectively.

The third storage partition consists of logical drives Marketing and HResources. This logical drive-to-LUN mapping is created using host group Omaha. These logical drives are accessed by host computers Omaha A and Omaha B in host group Omaha using LUNs 7 and 2, respectively.



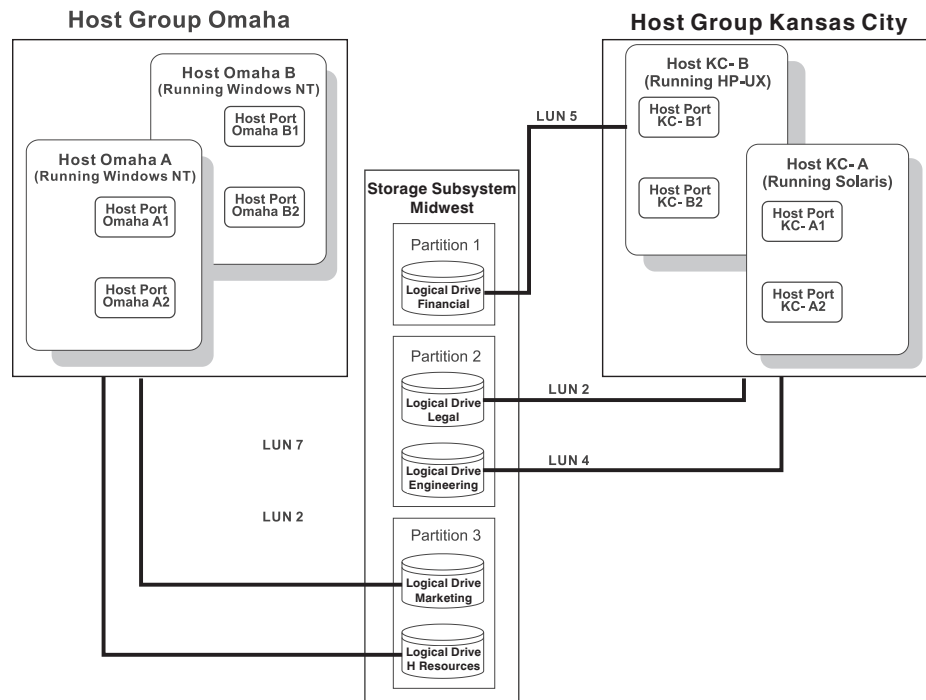


Figure 13. Storage-partitioning example

## Obtaining a feature key

Storage partitioning is enabled on a storage subsystem by default. However, if storage partitioning is not enabled, you must obtain a feature-key file. The feature-key file is created by sending your feature-enable identifier specific to the storage subsystem to your IBM Technical-Service Representative. You can obtain the feature-enable identifier by clicking **Storage subsystem** → **Premium Features** → **List** from the Subsystem Management window. After your IBM technical-service representative has returned a feature-key file, you can use it to enable storage partitioning by clicking **Storage subsystem** → **Premium Features** → **Enable** from the Subsystem Management window.

## Creating storage partitions

To create storage partitions, perform the following two tasks:

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

### Defining storage-partition topology

Use the **Configure** → **Storage Partitioning** option in the Subsystem Management window to display the Mappings window. Use the Mappings window to define the host groups, host computers, and associated host ports for which you want to define specific logical drive-to-LUN mappings (that is, the storage-partition topology).

The topological elements that you define (for example, host groups, host computers, and host ports) and the default host group are displayed in the Topology view. Figure 14 on page 30 shows the Topology view, which is located in the left

pane of the Mappings window.

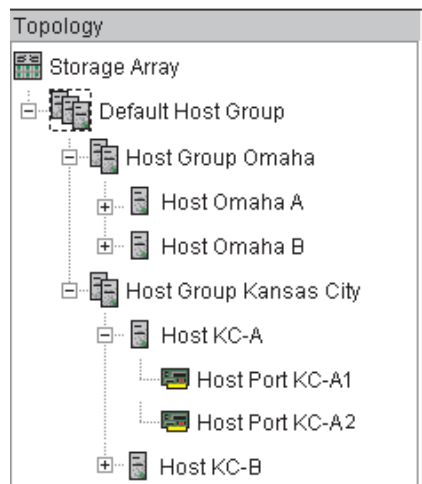


Figure 14. Topology View

Each host group or host computer that you define is granted a unique view of partitioned storage. A defined host group or host computer 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 computer is part of the default host group.

In Figure 14, the host groups Omaha and Kansas City and their associated host computers 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 host computers 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 computer or a host group can access either 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 in a defined host computer or a host group. To define specific logical drive-to-LUN mappings, complete the following steps:

1. Select a defined host computer or a host group in the Topology view.
2. Select the Define New Mapping option.

The Define New Mapping window opens, as shown in Figure 15 on page 31.

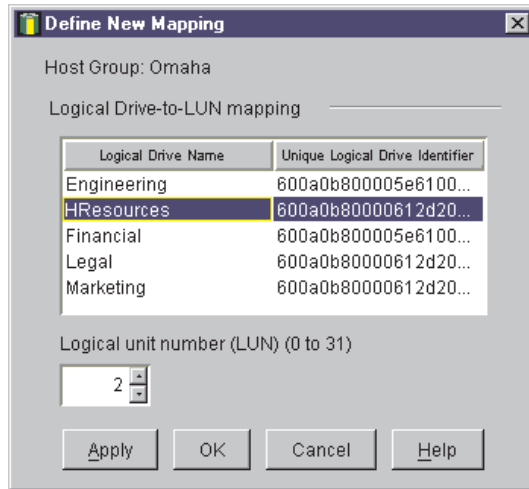


Figure 15. Define New Mapping window

3. In the Logical Drive-to-LUN mapping list, click a logical drive that you want the host group or host computer to access.
4. Select the logical unit number (LUN) that you want the host group or host computer to use to access the logical drive.

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

## Topology-view results

Figure 16 shows the results of creating logical drive-to-LUN mappings in the Topology view. After you define a specific logical drive-to-LUN mapping for host group Omaha, the host group is no longer a part of the default host group. The associated host computers (Omaha A and Omaha B) of host group Omaha now can access logical drive HResources, but because Omaha A and Omaha B are no longer part of the default host group, they cannot access any logical drives with default logical drive-to-LUN mappings.

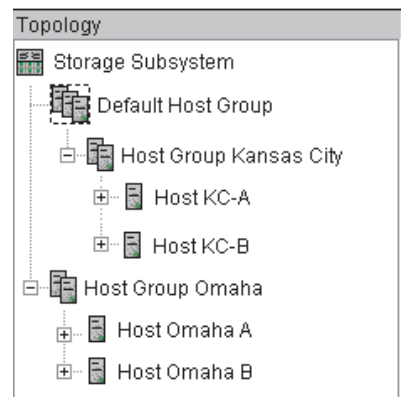
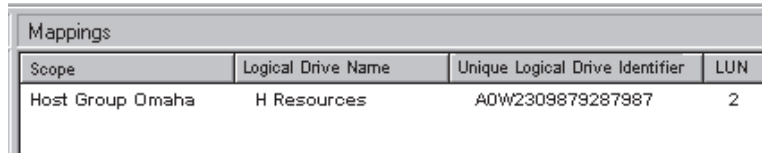


Figure 16. Creating a logical drive-to-LUN mapping in the Topology view

## Mappings-view results

After selecting the host group Omaha in the Topology view, the results appear in the Mappings view. Figure 17 shows the Mappings view, which is located in the right pane of the Mappings window.



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

Figure 17. Mappings view

For more information about creating logical drives or storage partitions, refer to the Subsystem Management window online help.

---

## Heterogeneous hosts overview

The Heterogeneous Hosts feature enables host computers running different operating systems to access a single storage subsystem. Previous releases of IBM FASiT Storage Manager enabled only host computers running the same operating system to access a single storage subsystem.

Host computers can be running different operating systems (for example, Sun Solaris and Windows NT) or variants of the same operating system (for example, Windows NT running in a cluster environment or Windows NT running in a noncluster environment). When a host-computer type is specified in the Define New Host Port window, the heterogeneous hosts feature enables the controllers in the storage subsystem to tailor their behavior (such as LUN reporting and error conditions) to the needs of the operating system or variant of the host computer sending the information. For detailed information about defining heterogeneous host computer types, see the Subsystem Management window online help.

**Important:** The Heterogeneous Hosts feature is available only when storage partitioning is enabled. In Figure 19 on page 33, storage partitioning is enabled.

In a heterogeneous environment, you must set each host-computer type to the appropriate operating system during host-port definition (Figure 18 on page 33). This enables firmware on each controller to respond correctly to the operating system of that host computer.

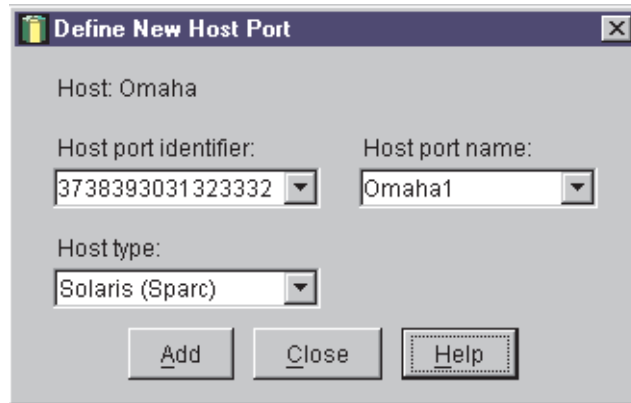


Figure 18. Define New Host Port window

In the example shown in Figure 19, both storage partitioning and heterogeneous host computers are enabled on storage subsystem Midwest. Four host computers (Omaha A and B and KC-A and B) are connected to storage subsystem Midwest. In the example, host computers Omaha A and Omaha B have access to logical drives Marketing and HResources.

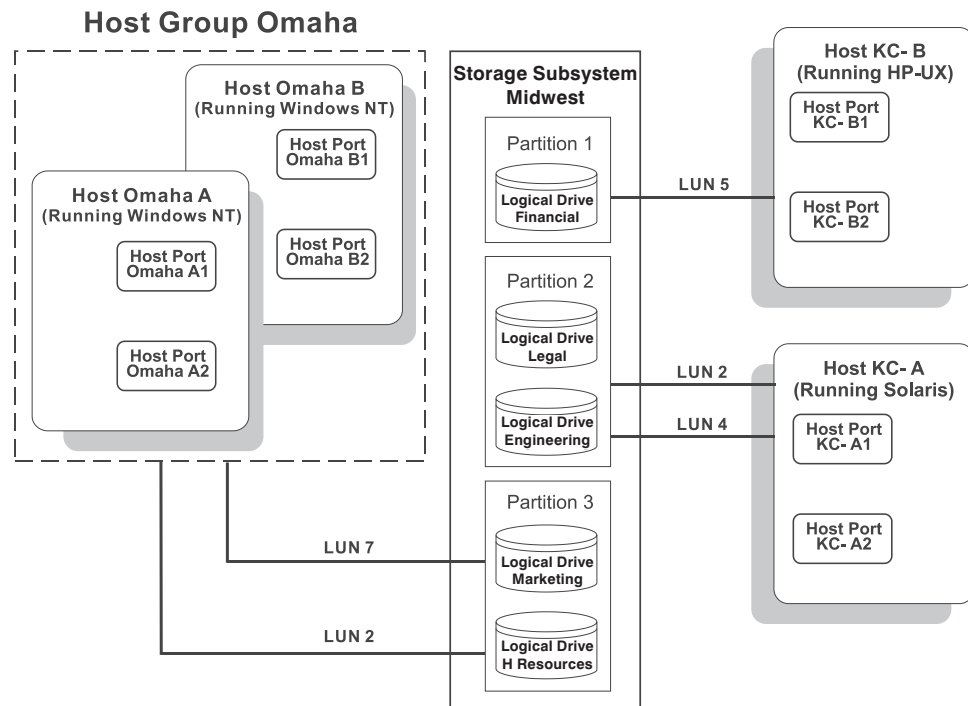


Figure 19. Heterogeneous hosts example

Host computer KC-A has exclusive access to logical drives Legal and Engineering, and host computer KC-B has exclusive access to logical drive Financial. Three storage partitions are configured, and because there are four host computers running three operating system environments, the heterogeneous hosts feature is needed to set the host-port types.

After you define the host-computer type for each host port, use one of the following two methods to display the host-computer type of the host port in the Topology view:

- Place your cursor over the specific host port and use the tooltip to display the associated host-computer type.
- Highlight the root, storage subsystem, and node. Click **Configure** → **Topology** → **Show All Host Port Information**.

## Chapter 4. Maintaining and monitoring storage subsystems

This chapter describes the tasks required for maintaining and monitoring storage subsystems in a management domain.

Use the Enterprise Management window to:

- Monitor the health status of the storage subsystems
- Configure alert destinations for critical event notification

**Note:** To receive critical alerts, the Enterprise Management window must be open (it can be minimized), or the Event Monitor must be installed and running.

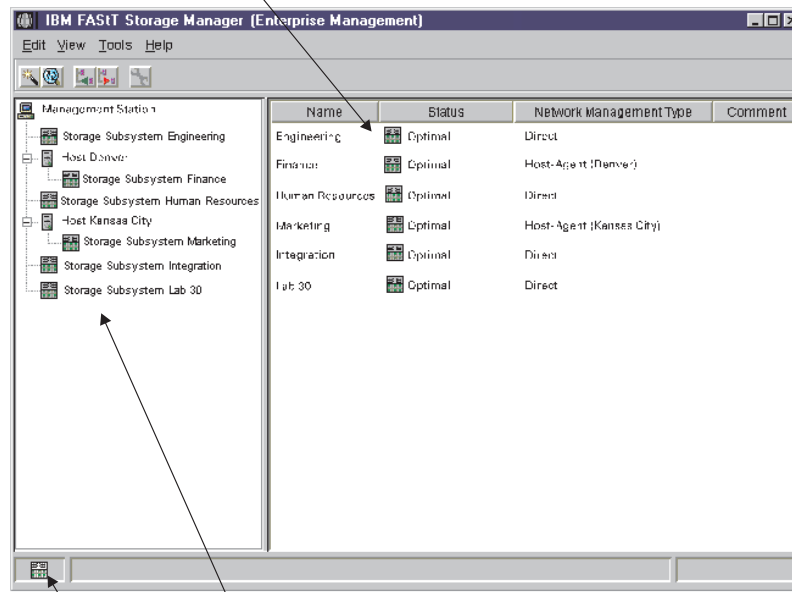
Use the Subsystem Management window to:

- Monitor the logical and physical components within a storage subsystem. See the Subsystem Management window online help.
- Monitor and tune storage-subsystem performance. See “Event Monitor overview” on page 39.
- Recover from storage-subsystem problems. See “Recovery Guru window” on page 41.

### Maintaining storage subsystems in a management domain

Use the Enterprise Management window to view the storage-subsystem status icons and monitor the storage-subsystem health.

The device table contains a health status column that displays the status of each storage subsystem.



The device tree in this pane provides the status for every storage subsystem in the management domain. A status is displayed 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.





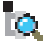
Figure 20. Monitoring storage-subsystem health using the Enterprise Management window

## Storage-subsystem status quick reference

Table 10 provides information about the storage-subsystem status icons that are displayed in the following areas:

- In the Device Tree, Device Table, and Overall Health Status panes of the Enterprise Management window
- As the root node of the Logical Tree view in the Subsystem Management window

Table 10. Storage-subsystem status icon quick reference

Status	Icon	Description
Optimal		An Optimal status indicates that every component in the storage subsystem is in the desired working condition.
Needs Attention		A Needs Attention status indicates that a problem on a storage subsystem requires intervention to correct. To correct the problem, open the Subsystem Management window for the particular storage subsystem; then, use the Recovery Guru to determine the cause of the problem and obtain the appropriate instructions.
Fixing		A Fixing status indicates that a Needs Attention condition has been corrected and the storage subsystem is going into an Optimal state (for example, a reconstruction operation is in progress). A Fixing status requires no action unless you want to check on the progress of the operation in the Subsystem Management window.  <b>Note:</b> Some recovery actions cause the storage-subsystem status to change directly from Needs Attention to Optimal, without an interim status of Fixing. In this case, the Fixing status icon is not displayed in the Overall Health Status pane. (Optimal is displayed instead.)
Unresponsive		An Unresponsive status indicates that the management station cannot communicate with the only controller or both controllers in the storage subsystem over its network-management connection.  <b>Note:</b> The Unresponsive icon is not displayed in the Logical view of the Subsystem Management window. If the Subsystem Management window is open and the storage subsystem becomes Unresponsive, the last known status (Optimal, Needs Attention, or Fixing) is displayed.
Contacting Device		A Contacting Device status indicates that you have opened the Enterprise Management window and the storage-management software is establishing contact with the storage subsystem.  <b>Note:</b> The Contacting Device status is not displayed in the Logical view of the Subsystem Management window.



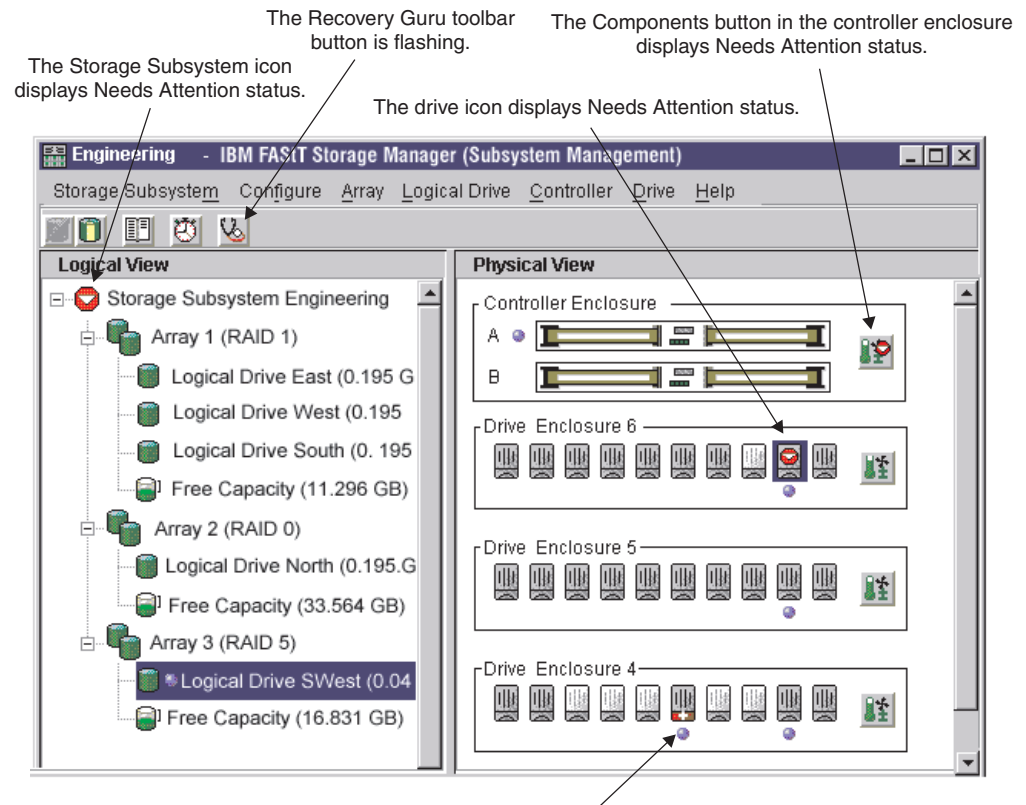
## Failure notification

When you monitor a storage subsystem, there are several indicators of a storage-subsystem failure. The Subsystem Management window displays the Needs Attention icon in the following locations:

- The Overall Health Status pane, Device Tree view, or Device Table of the Enterprise Management window.
- The Subsystem Management window Logical view.
- Individual storage subsystems.
- 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.
- Critical SNMP trap or e-mail error messages are sent.
- The hardware displays fault lights.

### Failure-notification example

Figure 21 shows the Subsystem Management window Logical View and the failure icons for storage subsystem Engineering.



The hot-spare drive is in use (taking over for a failed drive). The data from the failed drive has been reconstructed to the hot-spare drive. When the hot-spare drive has taken over for the failed drive, the associated logical drive (logical drive SWest) in logical drive group 3 remains in the Optimal status.

Figure 21. Failure notification in the Subsystem Management window

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

---

## Alert notification overview

This section provides a basis for understanding how alert notifications are sent. For detailed procedures, refer to the Enterprise Management window online help.

## Configuring mail server and sender address

To verify that the critical event information is sent, you must configure an e-mail server that forwards the e-mail to the configured e-mail alert destinations. Click **Edit** → **Configure Mail Server** in the Enterprise Management window. Next, you must specify the e-mail sender address (the address that will appear on every message). Typically, the e-mail sender address is the e-mail address of the network administrator.

## Selecting the node for notification

Alert notification settings are set at any level (management station, host computer, or storage subsystem). To choose to receive notifications for all storage subsystems monitored by a management station or a single storage subsystem on a host computer, select a node for notification in the Enterprise Management window. See the Enterprise Management window online help for more information.

## Setting alert destinations

You can choose to receive critical-event notifications through e-mail, SNMP traps, or both. Click **Edit** → **Alert Destinations** in the Enterprise Management window to type the destination information:

- On the e-mail Address page, type fully qualified e-mail addresses (name@company.com). Enter all the addresses where you want the information sent.
- On the SNMP traps page, type the community name and trap destination. The community name is set in the NMS by the network administrator. The trap destination is the IP address or the network-management station.

**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 *Storage Manager Installation Guide* for your operating system for details.

After alert destinations are set, a check mark is displayed in the left pane where the management station, host computer, or storage subsystem displays. When a critical problem occurs on the storage subsystem, the software sends a notification to the specified alert destinations.

## 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 host computer
- Individual storage subsystems

Also, you can use the storage-management software to validate potential destination addresses and specify management-domain global e-mail alert settings for mail-server and sender e-mail address.

---

## Event Monitor overview

The Event Monitor handles notification functions (e-mail and SNMP traps) and monitors storage subsystems whenever the Enterprise Management window is not open. Previous versions of the FAS*T* Storage Manager software did not have the Event Monitor and required that the Enterprise Management window remain open in order to monitor the storage subsystems and receive alerts. The Event Monitor is a separate program bundled with the Storage Manager 7 client software. (The Event Monitor cannot be installed without the client.) Install the client Event Monitor on a management station or host computer connected to the storage subsystems. For continuous monitoring, install the Event Monitor on a host computer that runs 24 hours a day. If you choose not to install the Event Monitor, you should still configure alerts on the host computer where the client software is installed.

The Event Monitor runs continuously in the background, monitoring activity on a storage subsystem and checks 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 through e-mail and SNMP. To configure alert notifications, click **Edit** → **Alert Destinations** in the Enterprise Management window. If the Event Monitor is configured and running on more than one host computer or management station connected to the storage subsystem, you will receive duplicate alert notifications for the same critical problem on that storage subsystem.

The Event Monitor and the FAS*T* Storage Manager 7 client software send alerts to a remote system. The `emwdata.bin` file on the management station contains the name of the storage subsystem being monitored and the address where to send alerts. The alerts and errors that occur on the storage subsystem are continually 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.

## Event-monitoring example

The Event Monitor and the Enterprise Management window share the information to send alert messages; therefore, the Enterprise Management window displays alert status to assist in Event Monitor installation and synchronization. The parts of the Enterprise Management window that are related to event monitoring are shown in Figure 22 on page 40.

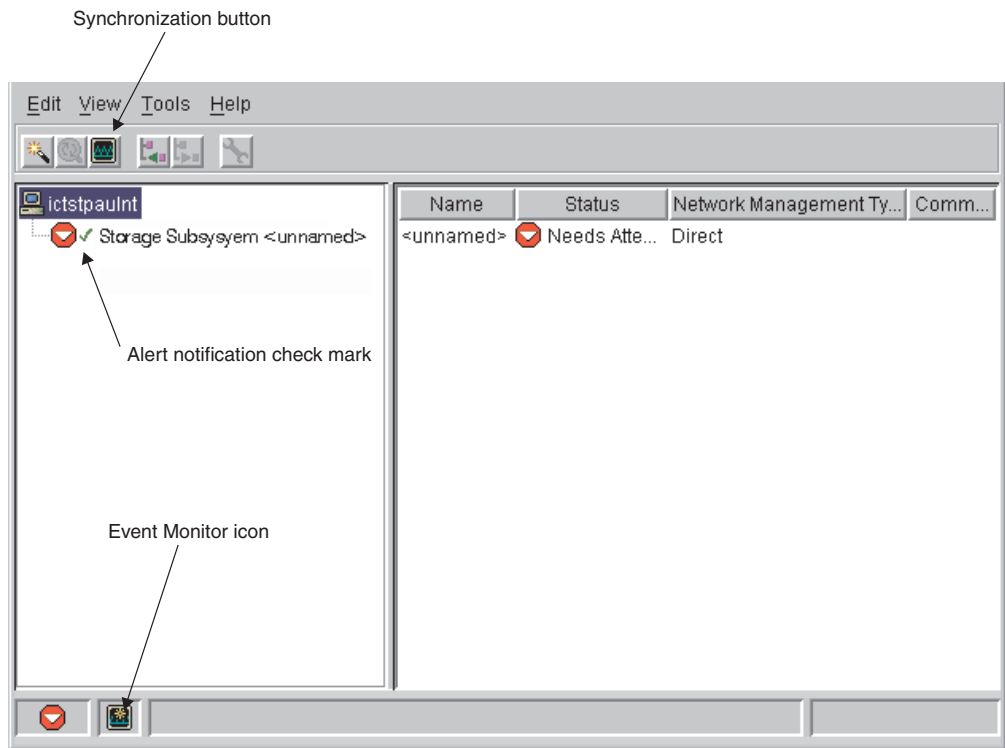


Figure 22. Event-monitoring example

## Installing the Event Monitor

The major steps in this section are provided as a basis for understanding event monitoring and alert notifications. For detailed procedures, see the Enterprise Management window online help. To install the Event Monitor software, you must have administrative permissions on the computer where the Event Monitor will reside, and you must install both the storage-manager client software and Event Monitor software together. After the software is installed, the Event Monitor icon (shown in Figure 22) is displayed in the lower-left corner of the Enterprise Management window.

## Setting alert notifications

To set up the alerts (e-mail and SNMP), click **Edit → Alert Destination** in the Enterprise Management window. A check mark indicates where the alert is set (management station, host computer, or storage subsystem). When a critical problem occurs on the storage subsystem, the Event Monitor sends a notification to the specified alert destinations.

The e-mail alert destinations will not work unless you also configure a mail server and sender e-mail address. Click **Edit → Configure Mail Server** in the Enterprise Management window. Configure the mail server and sender e-mail address only one time for all e-mail alert destinations.

**Note:** If you want to set identical alert destinations on more than one management station or host computer, you must install the Event Monitor on each system. Then you can either repeat setting up the alert destinations or copy the emwdata.bin file from one system to the other. However, be aware that if you have configured the Event Monitor on multiple systems that will monitor the

same storage subsystem, you will receive duplicate alert notifications for the same critical problem on that storage subsystem.

The emwdata.bin configuration file is stored in a default directory. The default directory is different, depending on your operating system. Locate and copy the file. After copying the file, remember to shut down and restart the Event Monitor and Enterprise Management window (or restart the host computer) for the changes to take effect. For more information, see the Enterprise Management window online help.

## Synchronizing the Enterprise Management window and Event Monitor

After the Event Monitor is installed, it continues to monitor storage subsystems and send alerts as long as it continues to run. When the Enterprise Management window is started, monitoring functions are shared by the Event Monitor and the Enterprise Management window. However, if you make a configuration change in the Enterprise Management window (such as adding or removing a storage subsystem or setting additional alert destinations), you must manually synchronize the Enterprise Management window and the Event Monitor using the **Synchronization** button, shown in Figure 22 on page 40.

When the Event Monitor and the Enterprise Management window are synchronized, the **Synchronization** button is unavailable. When a configuration change occurs, the Synchronization button becomes active. Clicking the button will synchronize the Event Monitor and the Enterprise Management software components.

**Note:** The Enterprise Management window and the Event Monitor are automatically synchronized whenever you exit from the Enterprise Management window. The Event Monitor continues to run and send alert notifications as long as the operating system is running.

For detailed information about setting up alert destinations or about the Enterprise Management window, see the Enterprise Management window online help.

---

## Recovery Guru window

The Recovery Guru is a component of the Subsystem Management window in the FAST Storage Manager client package. The Recovery Guru diagnoses storage-subsystem problems and suggests recovery procedures to correct the problems. To start the Recovery Guru, click the **Recovery Guru** toolbar button in the Subsystem Management window, shown in Figure 23 on page 42, or click **Storage Subsystem** → **Recovery Guru**.

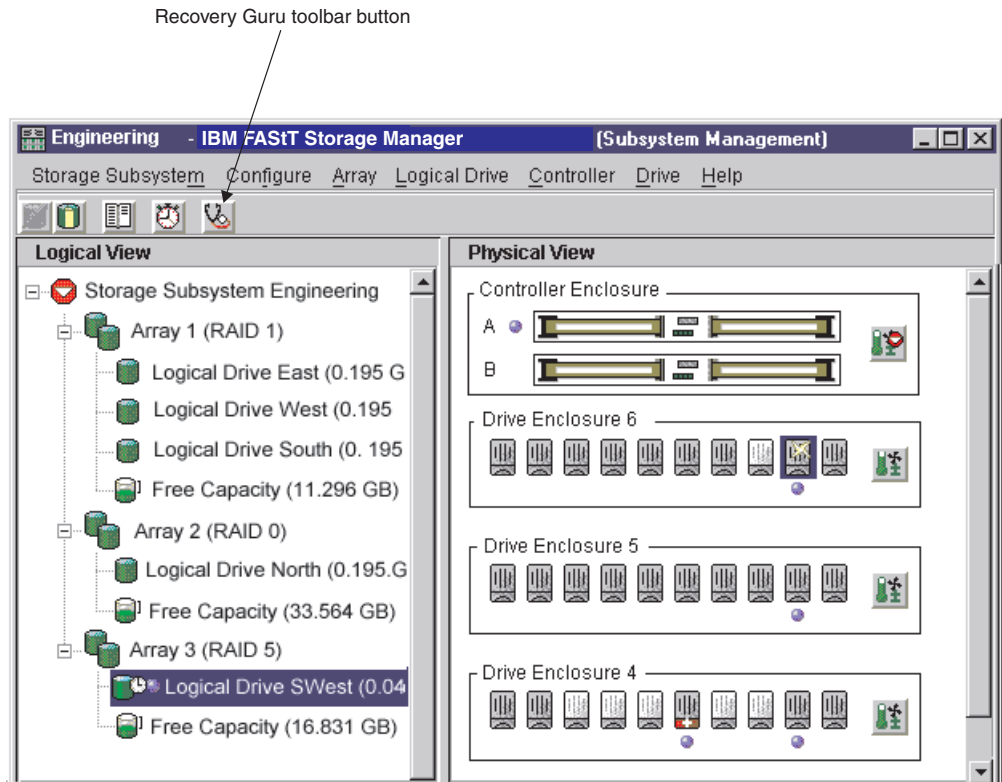


Figure 23. Location of the Recovery Guru toolbar button

## Recovering from storage-subsystem failures

The Recovery Guru window is shown in Figure 24 on page 43. The Summary pane shows that there are two different failures in this storage subsystem: a hot spare in use, and a failed battery customer replaceable unit (CRU).

When you select a failure from the list in the Summary pane, the appropriate details and a recovery procedure are displayed in the Details pane. For example, Figure 24 on page 43 shows that **Logical Drive - Hot Spare in Use** is selected. The Details pane shows that in logical drive SWest, a hot-spare drive has replaced a failed logical drive in enclosure 6, slot 9. The Recovery Procedure pane shows the details needed to recover from this failure.

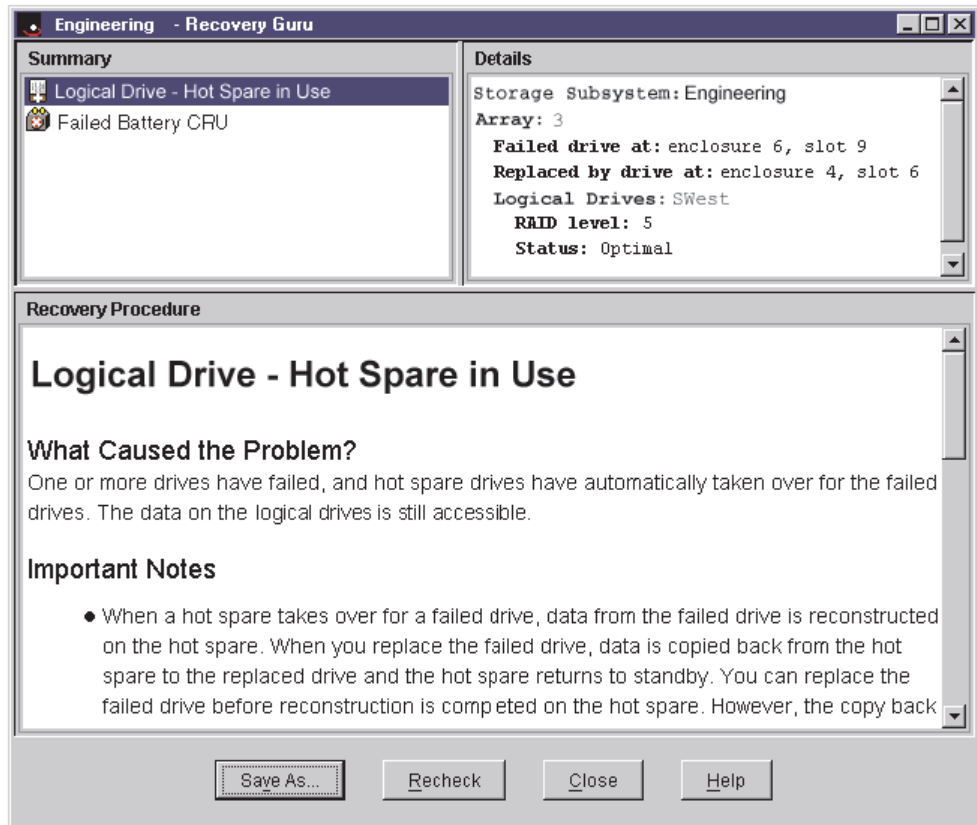


Figure 24. Recovery Guru window

As you follow the recovery procedure to replace the failed logical drive in the Subsystem Management window, the associated logical drive (SWest) icon changes to Operation in Progress, and the replaced logical drive icon changes to Replaced Drive status. The data that is reconstructed to the hot-spare drive is copied back to the replaced physical drive. During the copyback operation, the status icons change to reflect the Replaced status of the components, as shown in Figure 25 on page 44.

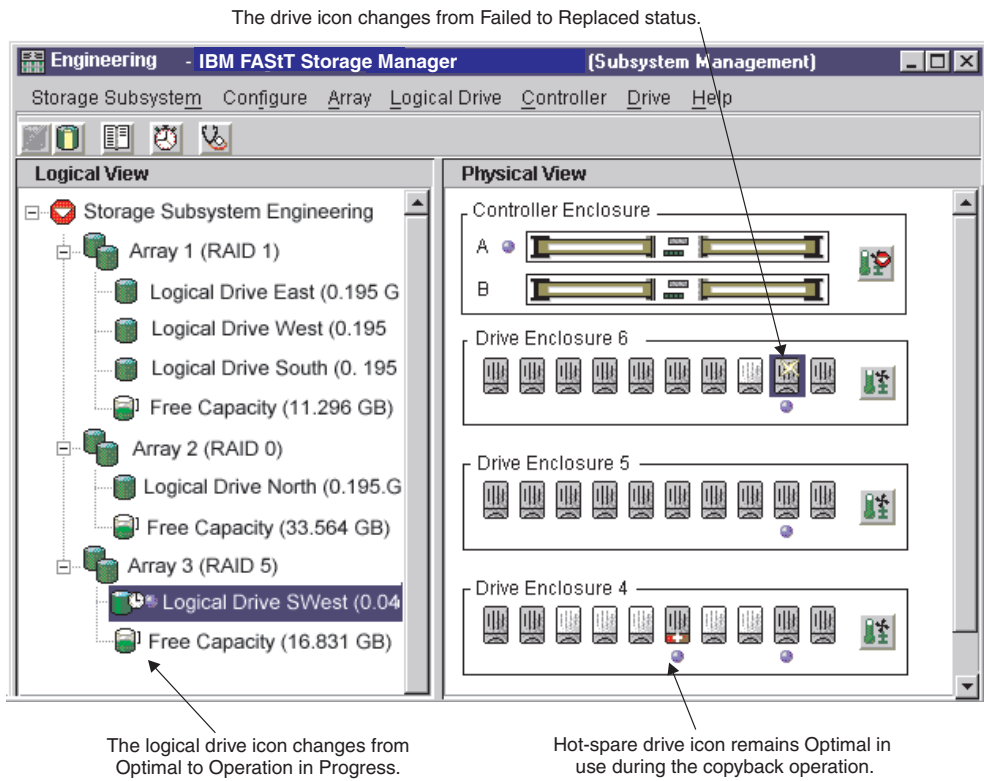
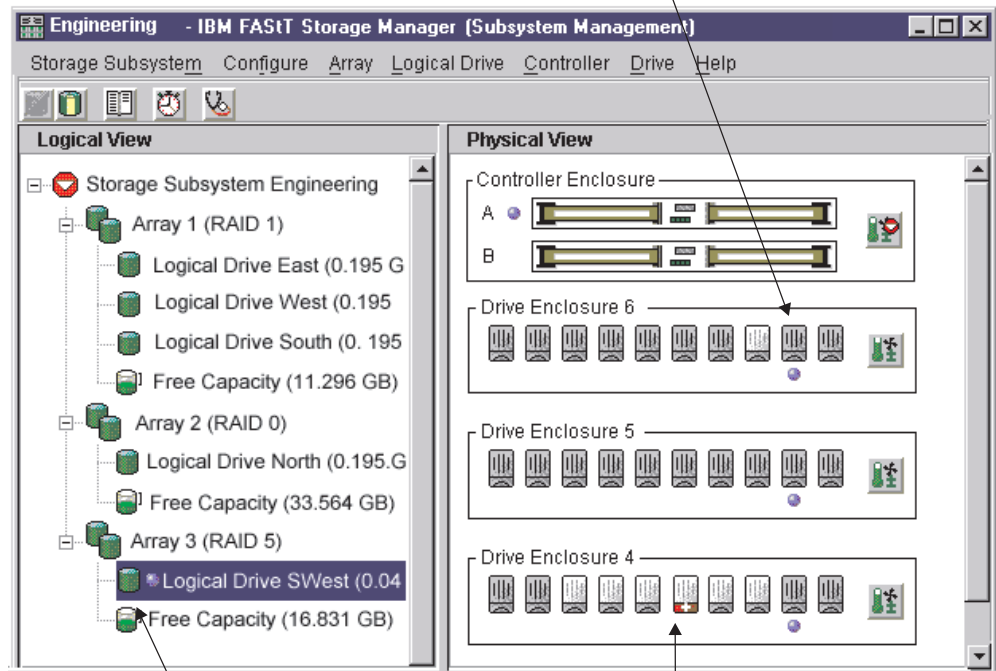


Figure 25. Recovery Guru window showing Replaced status

When the copyback operation is complete, the status icons change to reflect the Optimal status of the components, as shown in Figure 26 on page 45.



The drive icon changes from Replaced to Optimal Assigned.



The logical drive icon changes from Operation in Progress to Optimal.

The drive status icon for the hot-spare drive changes from hot-spare in use to hot-spare standby.

Figure 26. Recovery Guru window showing Optimal status

After you correct the storage subsystem errors:

- The Components button in the controller enclosure in the Physical view returns to Optimal.
- The storage-subsystem icon in the Logical view returns to Optimal.
- The storage-subsystem icon in the Enterprise Management window changes from Needs Attention to Optimal.
- The **Recovery Guru** toolbar button stops blinking.
- The components button in the controller enclosure status returns to Optimal, and the storage-subsystem status icon returns to Optimal.



---

## Chapter 5. Tuning storage subsystems

This section helps the system administrator use data from the Performance Monitor. This section also describes the tuning options available in IBM FAStT Storage Manager 7.10 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 11 describes the Performance Monitor data that is displayed for selected devices.

Table 11. Performance Monitor tuning options in the Subsystem Management window

Data field	Description
Total I/Os	Total I/Os performed by this device since the beginning of the polling session. For more information, see “Balancing the Fibre Channel I/O load”.
Read percentage	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 Fibre Channel I/O request rate” on page 48.
Cache-hit percentage	The percentage of read operations that are processed with data from the cache, rather than requiring a read from the logical drive. For more information, see “Optimizing the Fibre Channel I/O request rate” on page 48.
Current KB per second	During the polling interval, the <i>transfer rate</i> is the amount of data, in KB, that is moved through the Fibre Channel I/O path in one second (also called throughput). For more information, see “Optimizing the I/O transfer rate” on page 48.
Maximum KB per second	The maximum transfer rate that is achieved during the Performance Monitor polling session. For more information, see “Optimizing the I/O transfer rate” on page 48.
Current I/O per second	The average number of I/O requests serviced per second during the current polling interval (also called an I/O request rate). For more information, see “Optimizing the Fibre Channel I/O request rate” on page 48.
Maximum I/O per second	The maximum number of I/O requests serviced during a one-second interval over the entire polling session. For more information, see “Optimizing the Fibre Channel I/O request rate” on page 48.

---

### Balancing the Fibre Channel I/O load

The Total I/O data field in the Subsystem Management window is used for monitoring the Fibre Channel 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 Fibre Channel 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, move an array to the other controller in the storage subsystem by clicking **Array** → **Change Ownership**.

It is difficult to balance Fibre Channel 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 enables you to save data to a comma-delimited text file that you can import to a spreadsheet for further analysis.

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

## Optimizing the I/O transfer rate

The transfer rates of the controller are determined by the application I/O size and the I/O request rate. 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 will help you optimize the maximum I/O transfer rates that are possible for a given storage subsystem.

One of the ways to improve the I/O transfer rate is to improve the I/O request rate. Use the host-computer operating-system utilities to gather I/O size data to understand the maximum transfer rates possible. Then, use tuning options available in the IBM FAStT Storage Manager version 7.10 software to optimize the I/O request rate to reach the maximum possible transfer rate.

---

## Optimizing the Fibre Channel I/O request rate

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

- Fibre Channel 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 logical drives in the arrays or storage subsystem
- Fragmentation of files

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

- Logical-drive modification priority

## Determining the Fibre Channel I/O access pattern and I/O size

To determine if the Fibre Channel I/O access has sequential characteristics, try enabling a conservative cache read-ahead multiplier (for example, 4) by clicking **Logical Drive** → **Properties**. Then, examine the logical-drive cache-hit percentage to see if it has improved. An improvement indicates the Fibre Channel I/O has a sequential pattern. For more information, see “Optimizing the cache-hit percentage” on page 49. Use the host-computer operating-system utilities to determine the typical I/O size for a logical drive.

## Enabling write-caching

Higher Fibre Channel I/O write rates occur when write-caching is enabled, especially for sequential Fibre Channel I/O access patterns. Regardless of the Fibre Channel I/O access pattern, be sure to enable write-caching to maximize the Fibre Channel I/O rate and shorten application response time.

## Optimizing the cache-hit percentage

A higher cache-hit percentage is preferred for optimal application performance and is positively correlated with the Fibre Channel 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, you might need to install more memory.

If an individual logical drive has a low cache-hit percentage, consider enabling cache read-ahead (*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 fetches more data, usually from adjacent data blocks on the drive. In addition to the requested data, this feature increases the chance that a future request for data is fulfilled from the cache, rather than requiring a logical drive access.

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

If you determine that the Fibre Channel I/O access pattern has sequential characteristics, set an aggressive cache read-ahead multiplier (for example, 8). 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 0.)

## 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.

Applications with a low read percentage (write-intensive) do not perform as well on RAID level-5 logical drives because of the way a controller writes data and redundancy data to the drives in a RAID level-5 logical drive. If there is a low percentage of read activity relative to write activity, you can change the RAID level of a logical drive from RAID level-5 to RAID level-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 are completed, but the slower the system I/O access pattern 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 from the Logical Drive - Properties window. 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 Fibre Channel I/O request rate. If the controller is idle (not servicing any I/O request rates) 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 KB, that the controller writes on a single logical drive before writing data on the next drive. A data block is 512 bytes of data and is the smallest unit of storage. The size of a segment determines how many data blocks it contains. For example, an 8 KB segment holds 16 data blocks, and a 64 KB segment holds 128 data blocks.

**Important:** In IBM FASTT Storage Manager software versions 7.01 and 7.02, the segment size is expressed in the number of data blocks. The segment size in IBM Storage Manager version 7.10 is expressed in KB.

When you create a logical drive, the default-segment size is a good choice for the expected logical-drive usage. To change the default-segment size, click **Logical Drive** → **Change Segment Size**.

If the I/O size is larger than the segment size, increase the segment size 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 logical drive for a single request leaves other logical drives available to simultaneously service other requests.

When you use 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 logical drives in the array used for I/O requests). In this case, multiple logical drives are used for the same request, but each logical drive is accessed only 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. Verify that the files on the logical drive 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.

## Chapter 6. Task reference

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

*Table 12. Populating a management domain*

<b>Refer to the following online help</b>	<b>For information about the following 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 damaged configuration files
	Removing a device from a management domain

*Table 13. Configuring storage subsystems*

<b>Refer to the following online help</b>	<b>For information about the following 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 (defragmentation)
	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)
	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

Refer to the following online help	For information about the following task
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

Refer to the following online help	For information about the following task
Subsystem Management window	Changing a logical-drive LUN assignment, host-computer assignment, or host-group assignment
	Creating storage partitions
	Defining a logical drive-to-LUN mapping
	Deleting a host group, host computer, or host port from the defined storage-subsystem topology
	Deleting a logical drive-to-LUN mapping
	Granting logical-drive access to host computers
	Granting logical-drive access to host groups
	Moving a host computer from one host group to another host group
	Moving a host port from one host computer to another host computer
	Reconfiguring logical drive-to-LUN mappings
	Renaming a host group, host computer, 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

Refer to the following online help	For information about the following task
Subsystem Management window	Changing the RAID level of a logical drive
	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

Refer to the following online help	For information about the following task
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 computer
	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

Refer to the following online help	For information about the following 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

Refer to the following online help	For information about the following task
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

Refer to the following online help	For information about the following task
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

Refer to the following online help	For information about the following task
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



## Chapter 7. Critical event problem solving

When a critical event occurs, it is logged in the Event Log and also sent to any e-mail and SNMP trap destinations you have configured by clicking **Edit** → **Alert Destinations** in the Enterprise Management window. The critical event type and the sense key/ASC/ASCQ data are both shown in the event log details.

The following table provides more information about events with a critical priority, as shown in the Subsystem Management window event log.

Table 22. Critical events

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 1001 - Channel failed Sense key/ASC/ASCQ: 6/3F/C3	<p><b>Description:</b> The controller failed a channel and cannot access drives on this channel any more. The FRU group qualifier (byte 26) in the sense data indicates the 1-relative channel number of the failed channel. Typically this condition is caused by a drive ignoring the SCSI protocol on one of the controller destination channels. The controller fails a channel if it issued a reset on a channel and continues to see the drives ignore the SCSI Bus Reset on this channel.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed Drive SCSI Channel recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 1010 - Impending drive failure (PFA) detected Sense key/ASC/ASCQ: 6/5D/80	<p><b>Description:</b> A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive might fail within 24 hours.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Impending Drive Failure recovery procedure.</p>
Event 1015 - Incorrect mode parameters set on drive Sense key/ASC/ASCQ: 6/3F/BD	<p><b>Description:</b> The controller is unable to query the drive for its current critical mode page settings or is unable to change these settings to the correct setting. Currently, this indicates the Qerr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data.</p> <p><b>Action:</b> The controller has not failed yet. Contact your IBM technical-support representative for instructions on recovering from this critical event.</p>
Event 1207 - Fibre Channel link errors - threshold exceeded Sense key/ASC/ASCQ: None	<p><b>Description:</b> Invalid characters have been detected in the Fibre Channel signal. Possible causes for the error are a degraded laser in a Gigabit Interface converter (GBIC) or media interface adapter (MIA), damaged or faulty Fibre Channel cables, or poor cable connections between components on the loop.</p> <p><b>Action:</b> In the main Subsystem Management window, click <b>Help</b> → <b>Recovery Procedures</b>. Click <b>Fibre Channel Link Errors Threshold Exceeded</b> for more information about recovering from this event.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 150E - Controller loopback diagnostics failed Sense key/ASC/ASCQ: None	<p><b>Description:</b> The controller cannot initialize the drive-side Fibre Channel loops. A diagnostic routine has been run identifying a controller problem, and the controller has been placed offline. This event will occur only on certain controller models.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 202E - Read drive error during interrupted write Sense key/ASC/ASCQ: 3/11/8A	<p><b>Description:</b> A media error has occurred on a read operation during interrupted write processing.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 2109 - Controller cache not enabled - cache sizes do not match Sense key/ASC/ASCQ: 6/A1/00	<p><b>Description:</b> The controller cannot enable mirroring if the alternate controller cache size is different. Verify that both controllers have the same cache size.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 210C - Controller cache battery failed Sense key/ASC/ASCQ: 6/0C/80	<p><b>Description:</b> The controller has detected that the battery is not physically present, is fully discharged, or has reached its expiration date.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed Battery CRU recovery procedure.</p>
Event 210E - Controller cache memory recovery failed after power cycle or reset Sense key/ASC/ASCQ: 6/0C/81	<p><b>Description:</b> Recovery from a data-cache error is unsuccessful. User data might have been lost.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 2110 - Controller cache memory initialization failed Sense key/ASC/ASCQ: 6/40/81	<p><b>Description:</b> The controller has detected the failure of an internal controller component (RAID buffer). The internal controller component failure might have been detected during operation as well as during an on-board diagnostic routine.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 2113 - Controller cache battery nearing expiration Sense key/ASC/ASCQ: 6/3F/D9	<p><b>Description:</b> The cache battery expiration is within the specified number of weeks.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Battery Nearing Expiration recovery procedure.</p>
Event 2229 - Drive failed by controller Sense key/ASC/ASCQ: None	<p><b>Description:</b> The controller has failed a drive because of a problem with the drive.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the procedures to correct the failure.</p>
Event 222D - Drive manually failed Sense key/ASC/ASCQ: 6/3F/87	<p><b>Description:</b> The drive is manually failed by a user.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 2247 - Data lost on the the logical drive during unrecovered interrupted write Sense key/ASC/ASCQ: 6/3F/EB	<p><b>Description:</b> An error has occurred during interrupted write processing during the start-of-day routine, causing the logical drive to go into a failed state.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 2248 - Drive failed - write failure Sense key/ASC/ASCQ: 6/3F/80	<p><b>Description:</b> The drive failed during a write command. The drive will be marked failed.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 2249 - Drive capacity less than minimum Sense key/ASC/ASCQ: 6/3F/8B	<p><b>Description:</b> During drive replacement, the capacity of the new drive is not large enough to support all the logical drives that must be reconstructed on it.</p> <p><b>Action:</b> Replace the drive with a larger-capacity drive.</p>
Event 224A - Drive has wrong block size Sense key/ASC/ASCQ: 6/3F/8C	<p><b>Description:</b> The drive block size does not match that of the other drives in the logical drive. The drive will be marked failed.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 224B - Drive failed - initialization failure Sense key/ASC/ASCQ: 6/3F/86	<p><b>Description:</b> The drive failed either from a Format Unit command, or a Write operation (issued when a logical drive is initialized). The drive will be marked failed.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 224D - Drive failed - no response at start of day Sense key/ASC/ASCQ: 6/3F/85	<p><b>Description:</b> The drive failed a Read Capacity or Read command during the start-of-day routine. The controller is unable to read the configuration information stored on the drive. The drive will be marked failed.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 224E - Drive failed - initialization/reconstruction failure Sense key/ASC/ASCQ: 6/3F/82	<p><b>Description:</b> The (previously-failed) drive is marked failed either because the drive failed a Format Unit command issued to it or because the reconstruction on the drive failed due to the controller being unable to restore it (for example, an error occurring on another drive required for reconstruction).</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 2250 - Logical drive failure (3F E0) Sense key/ASC/ASCQ: 6/3F/E0	<p><b>Description:</b> The controller has marked the logical drive failed. User data and redundancy (parity) can no longer be maintained to ensure availability. The most likely cause is the failure of a single drive in nonredundant configurations or a nonredundant second drive in a configuration protected by one drive.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed Logical Drive Failure recovery procedure.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 2251 - Drive failed - reconstruction failure Sense key/ASC/ASCQ: 6/3F/8E	<p><b>Description:</b> A drive failed due to a reconstruction failure during the start-of-day routine.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 2252 - Drive marked offline during interrupted write Sense key/ASC/ASCQ: 6/3F/98	<p><b>Description:</b> An error has occurred during interrupted write processing, causing the logical drive to be marked failed. Drives in the array that did not experience the read error will go into the Offline state and log this error.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 2254 - Redundancy (parity) and data mismatch is detected Sense key/ASC/ASCQ: 6/8E/01	<p><b>Description:</b> The controller detected inconsistent redundancy (parity) or data during a parity verification.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 2255 - Logical drive definition incompatible with ALT mode - ALT disabled Sense key/ASC/ASCQ: 6/91/3B	<p><b>Description:</b> Auto-LUN Transfer (ALT) works only with arrays that have only one logical drive defined. There are currently arrays on the storage subsystem that have more than one logical drive defined; therefore, ALT mode has been disabled. The controller will operate in normal redundant controller mode, and if there is a problem, it will transfer all logical drives on an array instead of transferring individual logical drives.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 2602 - Automatic controller firmware synchronization failed Sense key/ASC/ASCQ: 02/04/81	<p><b>Description:</b> The versions of firmware on the redundant controllers are not the same because the automatic controller firmware synchronization failed. Controllers with an incompatible version of firmware might cause unexpected results.</p> <p><b>Action:</b> Try the firmware download again. If the problem persists, contact your IBM technical-support representative.</p>
Event 2801 - Storage subsystem running on UPS battery Sense key/ASC/ASCQ: 6/3F/C8	<p><b>Description:</b> The uninterruptible power supply (UPS) has indicated that ac power is no longer present and the UPS has switched to standby power. While there is no immediate cause for concern, you should save your data frequently, in case the battery is suddenly depleted.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Lost ac Power recovery procedure.</p>
Event 2803 - UPS battery - two minutes to failure Sense key/ASC/ASCQ: 6/3F/C9	<p><b>Description:</b> The uninterruptible power supply (UPS) has indicated that its standby power supply is nearing depletion.</p> <p><b>Action:</b> Take actions to stop I/O activity to the controller. Normally, the controller will change from a write-back caching mode to a write-through mode.</p>
Event 2804 - UPS battery failed Sense key/ASC/ASCQ: None	<p><b>Description:</b> The UPS battery has failed.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>



Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 2807 - Environmental service monitor failed Sense key/ASC/ASCQ: None	<p><b>Description:</b> An environmental service monitor has failed.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed Environmental service monitor CRU recovery procedure.</p>
Event 2808 - Enclosure ID not unique Sense key/ASC/ASCQ: 6/98/01	<p><b>Description:</b> The controller has determined that there are multiple drive enclosures with the same ID selected. Verify that each drive enclosure has a unique ID setting.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Enclosure ID Conflict recovery procedure.</p>
Event 280A - Controller enclosure component missing Sense key/ASC/ASCQ: 6/3F/C7	<p><b>Description:</b> A component other than a controller is missing in the controller enclosure (for example, a fan, power supply, or battery). The FRU codes indicate the faulty component.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 280B - Controller enclosure component failed Sense key/ASC/ASCQ: 6/3F/C7	<p><b>Description:</b> A component other than a controller has failed in the controller enclosure (for example, a fan, power supply, battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component.</p> <p><b>Action:</b> Select Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 280D - Drive enclosure component failed Sense key/ASC/ASCQ: 6/3F/C7	<p><b>Description:</b> A component other than a drive has failed in the drive enclosure (for example, a fan, power supply, or battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the recovery procedure and follow the instructions to correct the failure.</p>
Event 280E - Standby power supply not fully charged Sense key/ASC/ASCQ: 6/3F/CA	<p><b>Description:</b> The uninterruptible power supply (UPS) has indicated that its standby power supply is not at full capacity.</p>
Event 280F - Environmental service monitor - loss of communication Sense key/ASC/ASCQ: 6/E0/20	<p><b>Description:</b> Communication has been lost to one of the dual environmental service monitor CRUs in a drive enclosure. The drive enclosure has only one I/O path available.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Environmental service monitor CRU - Loss of Communication recovery procedure.</p>
Event 2813 - Mini-hub CRU failed Sense key/ASC/ASCQ: 6/3F/C7	<p><b>Description:</b> Communication with the mini-hub CRU has been lost. This might be the result of a mini-hub CRU failure, a controller failure, or a failure in an internal backplane communications board. If there is only one mini-hub failure, the storage subsystem is still operational, but a second mini-hub failure could result in the failure of the affected enclosures.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed Mini-Hub CRU recovery procedure.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 2815 - GBIC failed Sense key/ASC/ASCQ: None	<p><b>Description:</b> A Gigabit Interface Converter (GBIC) on either the controller enclosure or the drive enclosure has failed. If there is only one GBIC failure, the storage subsystem is still operational, but a second GBIC failure could result in the failure of the affected enclosures.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Failed GBIC recovery procedure.</p>
Event 2816 - Enclosure ID Conflict - duplicate IDs across drive enclosures Sense key/ASC/ASCQ: 6/98/01	<p><b>Description:</b> Two or more drive enclosures are using the same enclosure identification number.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Enclosure ID Conflict recovery procedure.</p>
Event 2818 - Enclosure ID mismatch - duplicate IDs in the same drive enclosure Sense key/ASC/ASCQ: 6/98/02	<p><b>Description:</b> A drive enclosure in the storage subsystem contains environmental service monitors with different enclosure identification numbers.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Enclosure ID Mismatch recovery procedure.</p>
Event 281B - Nominal temperature exceeded Sense key/ASC/ASCQ: 6/98/03	<p><b>Description:</b> The nominal temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. If the temperature of the enclosure continues to rise, the affected enclosure might automatically shut down. Fix the problem immediately, before it becomes more serious. The automatic shutdown conditions depend on the model of the enclosure.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Nominal Temperature Exceeded recovery procedure.</p>
Event 281C- Maximum temperature exceeded Sense key/ASC/ASCQ: 6/3F/C6	<p><b>Description:</b> The maximum temperature of the enclosure has been exceeded. Either a fan has failed or the temperature of the room is too high. This condition is critical and might cause the enclosure to shut down if you do not fix the problem immediately. The automatic shutdown conditions depend on the model of the enclosure.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Maximum Temperature Exceeded recovery procedure.</p>
Event 281D - Temperature sensor removed Sense key/ASC/ASCQ: 6/98/03	<p><b>Description:</b> A CRU containing a temperature sensor has been removed from the storage subsystem.</p> <p><b>Action:</b> Replace the CRU as soon as possible. Start the Recovery Guru to obtain the Failed or Removed Fan CRU recovery procedure.</p>
Event 281E - Environmental service monitor firmware mismatch Sense key/ASC/ASCQ: 6/98/03	<p><b>Description:</b> A drive enclosure in the storage subsystem contains environmental service monitors with different versions of firmware. environmental service monitors in the same drive enclosure must have the same version firmware. If you do not have a replacement service monitor, call your IBM technical-support representative to perform the firmware download.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Environmental service monitor Firmware Version Mismatch recovery procedure.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 3019 - Logical drive ownership changed due to failover Sense key/ASC/ASCQ: None	<p><b>Description:</b> The multipath driver software has changed ownership of the logical drives to the other controller because it could not access the logical drives on the particular path.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Logical Drive Not on Preferred Path recovery procedure.</p>
Event 5005 - Place controller offline Sense key/ASC/ASCQ: None	<p><b>Description:</b> The controller is placed offline. This could be caused by the controller failing a diagnostic test. (The diagnostics are initiated internally by the controller or by the <b>Controller</b> → <b>Run Diagnostics</b> menu option.) Or the controller is manually placed Offline using the <b>Controller</b> → <b>Place Offline</b> menu option.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 5602 - This controller's alternate failed - timeout waiting for results Sense key/ASC/ASCQ: None	<p><b>Description:</b> This controller initiated diagnostics on the alternate controller but did not receive a reply indicating that the diagnostics were completed. The alternate controller in this pair has been placed offline.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 560B - CtlrDiag task cannot obtain Mode Select lock Sense key/ASC/ASCQ: None	<p><b>Description:</b> This controller is attempting to run diagnostics and could not secure the test area from other storage-subsystem operations. The diagnostics were canceled.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 560C - CtlrDiag task on controller's alternate cannot obtain Mode Select lock Sense key/ASC/ASCQ: None	<p><b>Description:</b> The alternate controller in this pair is attempting to run diagnostics and could not secure the test area from other storage-subsystem operations. The diagnostics were canceled.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 560D - Diagnostics read test failed on controller Sense key/ASC/ASCQ: None	<p><b>Description:</b> While running diagnostics, the controller detected that the information received does not match the expected return for the test. This could indicate that I/O is not completing or that there is a mismatch in the data being read. The controller is placed offline as a result of this failure.</p> <p><b>Action:</b> Select the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>

Table 22. Critical events (continued)

Critical event type and sense key/ASC/ASCQ	Critical event description and required action
Event 560E - This controller's alternate failed diagnostics read test Sense key/ASC/ASCQ: None	<p><b>Description:</b> While running diagnostics, the alternate for this controller detected that the information received does not match the expected return for the test. This could indicate that I/O is not completing or that there is a mismatch in the data being read. The alternate controller in this pair is placed offline.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 560F - Diagnostics write test failed on controller Sense key/ASC/ASCQ: None	<p><b>Description:</b> While running diagnostics, the alternate for this controller is unable to write data to the test area. This could indicate that I/O is not being completed or that there is a mismatch in the data being written. The controller is placed offline.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 5610 - This controller's alternate failed diagnostics write test Sense key/ASC/ASCQ: None	<p><b>Description:</b> While running diagnostics, the alternate for this controller is unable to write data to the test area. This could indicate that I/O is not being completed or that there is a mismatch in the data being written. The alternate controller in this pair is placed offline.</p> <p><b>Action:</b> Start the Recovery Guru to obtain the Offline Controller recovery procedure. Use this procedure to replace the controller.</p>
Event 5616 - Diagnostics rejected - configuration error on controller Sense key/ASC/ASCQ: None	<p><b>Description:</b> The alternate for this controller is attempting to run diagnostics and could not create the test area necessary for the completion of the tests. The diagnostics were canceled.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>
Event 5617 - Diagnostics rejected - configuration error on controller's alternate Sense key/ASC/ASCQ: None	<p><b>Description:</b> The alternate for this controller is attempting to run diagnostics and could not create the test area necessary for the completion of the tests. The diagnostics were canceled.</p> <p><b>Action:</b> Contact your IBM technical-support representative for instructions on recovering from this failure.</p>

---

## Appendix A. Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your computer, and whom to call for service should it be necessary.

---

### Getting information

Information about your IBM server product and preinstalled software, if any, is available in the documentation that comes with your server. That documentation includes printed books, online books, README files, and help files. In addition, information about IBM products is available on the World Wide Web and through the IBM Automated Fax System.

#### Using the World Wide Web

On the World Wide Web, the IBM Web site has up-to-date information about IBM products and support. The address for the IBM Personal Computing home page is <http://www.ibm.com/pc/>.

You can find support information for your IBM products at <http://www.ibm.com/pc/support/>.

If you click **Profile** from the support page, you can create a customized support page that is specific to your hardware, complete with Frequently Asked Questions, Parts Information, Technical Hints and Tips, and Downloadable Files. In addition, you can choose to receive e-mail notifications whenever new information becomes available about your registered products.

You also can order publications through the IBM Publications Ordering System at <http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>.

#### Getting information by fax

If you have a touch-tone telephone and access to a fax machine, in the U.S. and Canada, you can receive, by fax, marketing and technical information on many topics, including hardware, operating systems, and local area networks (LANs).

You can call the IBM Automated Fax System 24 hours a day, 7 days a week. Follow the recorded instructions, and the requested information will be sent to your fax machine. In the U.S. and Canada, to access the IBM Automated Fax System, call 1-800-426-3395.

---

### Getting help and service

If you have a problem with your server product, you will find a wide variety of sources available to help you.

#### Using the documentation and diagnostic programs

Many problems are solved without outside assistance. If you experience a problem with your server product, the first place to start is the troubleshooting information in

your IBM documentation. If you suspect a software problem, see the documentation, including README files and online help, that comes with the operating system or application program.

Most IBM server products come with a set of diagnostic programs that you can use to help you identify hardware problems. See the troubleshooting information in your IBM documentation for instructions on using the diagnostic programs.

The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/pc/support/> and follow the instructions.

## Calling for service

If you have tried to correct the problem yourself and still need help, during the warranty period, you can get help and information by telephone through the IBM HelpCenter«. The following services are available during the warranty period:

- **Problem determination** - Trained personnel are available to assist you with determining if you have a hardware problem and deciding what action is necessary to fix the problem.
- **IBM hardware repair** - If the problem is determined to be caused by IBM hardware under warranty, trained service personnel are available to provide the applicable level of service.
- **Engineering Change management** - Occasionally, there might be changes that are required after a product has been sold. IBM or your reseller, if authorized by IBM, will make selected Engineering Changes (ECs) available that apply to your hardware.

The following items are not covered:

- Replacement or use of non-IBM parts or nonwarranted IBM parts. All warranted parts contain a 7-character identification in the format IBM FRU XXXXXXX.
- Identification of software problem sources.
- Configuration of BIOS as part of an installation or upgrade.
- Changes, modifications, or upgrades to device drivers.
- Installation and maintenance of network operating systems (NOS).
- Installation and maintenance of application programs.

Refer to your IBM Statement of Limited Warranty for a full explanation of IBM warranty terms. Be sure to retain your proof of purchase to obtain warranty service.

In the U.S. and Canada, these services are available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9:00 a.m. to 6:00 p.m.

**Note:** Response times vary; may exclude some holidays.

In addition, you are eligible for IBM Start Up Support for 90 days after installation. This service provides assistance for:

- Setting up your network operating system
- Installing and configuring interface adapters
- Installing and configuring network adapters

Please have the following information ready when you call:

- Machine type and model

- Serial numbers of your IBM hardware products
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

## Telephone numbers

Phone numbers are subject to change without notice. For the most up-to-date phone numbers, go to <http://www.ibm.com/pc/support/> and click **Support Phone List**.

Country		Telephone number
Austria	Österreich	01-24 592 5901
Belgium - Dutch	Belgie	02-210 9820
Belgium - French	Belgique	02-210 9800
Canada	Toronto only	416-383-3344
Canada	Canada - all other	1-800-565-3344
Denmark	Danmark	45 20 82 00
Finland	Suomi	09-22 931 840
France	France	02 38 55 74 50
Germany	Deutschland	07032-1549 201
Ireland	Ireland	01-815 9202
Italy	Italia	02-7031 6101
Luxembourg	Luxembourg	298-977 5063
Netherlands	Nederland	020-514 5770
Norway	Norge	66 81 11 00
Portugal	Portugal	21-791 51 47
Spain	España	91-662 49 16
Sweden	Sverige	08-477 4420
Switzerland	Schweiz/Suisse/Svizzera	058-333 0900
United Kingdom	United Kingdom	01475-555 055
U.S.A. and Puerto Rico	U.S.A. and Puerto Rico	1-800-772-2227

In all other countries or regions, contact your IBM reseller or IBM marketing representative.

---

## International Warranty Service

If you travel with your server product or relocate it to a country where your server machine type is sold and serviced by IBM or by IBM resellers who are authorized to perform warranty service, International Warranty Service is available for your server product.

Service methods and procedures vary by country, and some services might not be available in all countries. International Warranty Service is delivered through the method of service (such as depot, carry-in, or on-site service) that is provided in the servicing country. Service centers in certain countries might not be able to service all models of a particular machine type. In some countries, fees and restrictions might apply at the time of service.

To determine whether your server product is eligible for International Warranty Service and to view a list of the countries where service is available, go to <http://www.ibm.com/pc/support/>, and click **Warranty Lookup**. Eligible IBM server products are identified by their four-digit machine types. For more information about International Warranty Service, see Service Announcement 601-034 at <http://www.ibm.com/>. Click **Support & downloads**, click **Announcement letters**, and click **Announcements search**. In the **Search by letter number only** field, type 601-034, and click **Search**.

---

## Purchasing additional services

During and after the warranty period, you can purchase additional services, such as support for IBM and non-IBM hardware, operating systems, and application programs; network setup and configuration; upgraded or extended hardware repair services; and custom installations. Service availability and service name might vary by country or region.

For more information about these services, contact your IBM marketing representative.



---

## Appendix B. Notices

This publication was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. 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. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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.*

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 NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information 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 without notice.

Any references in this publication to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

---

### Edition notice

**© COPYRIGHT INTERNATIONAL BUSINESS MACHINES CORPORATION, 1999, 2001. 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.

---

## Trademarks

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

AIX

e-business logo

HelpCenter

IBM

Netfinity

ServerProven

xSeries

Intel, Celeron, MMX, LANDesk, Pentium, Pentium II Xeon, Pentium III Xeon, and Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

---

## Important notes

Processor speeds indicate the internal clock speed of the microprocessor; other factors also affect application performance.

CD-ROM drive speeds list the variable read rate. Actual speeds vary and are often less than the maximum possible.

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1000000 bytes, and GB stands for approximately 1000000000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Maximum memory may require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven«, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

Unless otherwise stated, IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.



---

# Index

## A

- access volume 8
- add device option 9
- alert destinations 35
  - configuration 38
  - setting 38
- alert notification
  - setting 40
- alert notification overview 38
- automatic discovery option 9

## C

- cache flush
  - described 20
  - performance impacts 21
  - settings 21
  - start percentage 20
  - stop flush percentage 21
- cache hit
  - optimizing 49
- cache hit percentage 49
- cache read-ahead, choosing a multiplier
  - defined 48
- channel protection 22
- components button 45
- Contacting Device status 36
- controller enclosure 21
- controller transfer rate optimizing 48
- copyback 21
- creating logical drives from 25
- critical event notification 38
  - through SNMP 38
- customer support alert notification
  - how to configure 38

## D

- data protection strategies 19
- default host group, defined 27, 30
- default logical drive-to-LUN mapping
  - defined 27
- default LUN 27
- defragmenting files 50
- Device Table 36
- Device Tree 36
- DHCP/BOOTP 12
- directly managed storage subsystems 5

## E

- elements, storage topology 25
- emwdata.bin file 16
- enable a background media scan 21
- Enterprise Management window
  - component of SM7client 14
  - help 1

- Enterprise Management window (*continued*)
  - monitoring storage subsystems 35
  - status icons displayed in 36
- event log 53, 57
- Event Monitor
  - installing 40
  - overview 39
- event notification, configuring alert destinations 38
- existing storage subsystem, system hardware requirements 11

## F

- fabric switches 11
- failure notification 37
  - in the Subsystem Management window 37
- Fibre Channel Concepts Guide described 1
- Fibre Channel I/O
  - access pattern 48
  - balancing the load 47
  - request rate 48
  - request rate optimizing 48
  - size 48
- firmware requirements 13
- Fixing status 36
- free capacity node
  - adding to a logical drive group 25

## H

- hardware components
  - DHCP server, BOOTP or BOOTP compliant 11
  - host computer 12
  - management station 11
  - network-management station 11
  - storage subsystem 12
- hardware requirements 11
- heterogeneous hosts
  - defining types 32
  - example 33
  - overview 32
- host adapters 11
- host computer defined as topological element 27
- host group
  - defined 27
- host port
  - defined 26, 32
  - discovery of 26
- host type 34
- host-agent managed method 10
- host-agent management method
  - advantages 7
  - described 7
  - disadvantages 8
- hot spare drive
  - configuring 21
  - defined 21

## I

- I/O data field 47
- I/O data path protection 22
- I/O request rate, impact from cache flush settings 21

## L

- logical drive
  - creating 23
  - expected usage 25
  - modification priority setting 49
- logical drive creation, step-by-step 23
- logical drive group, capacity of 19
- logical drive-to-LUN mapping
  - default 27
  - defined 27, 30
  - from the Mappings view 32
  - specific 27
- logical drive-to-LUN terminology
  - default host group 27
  - host 27
  - host group 27
  - host port 26
  - mapping 27
  - storage partition 26
  - storage partition topology 26
  - storage partitions
    - mapping preference 28
- logical unit number (LUN) 8
- Logical view 36
- logical view, subsystem-management window 17
- LUN
  - address space 27
  - defined 27

## M

- mail server configuration 38
- managed hub 11
- management methods for storage subsystem
  - direct 5
  - host-agent 5
- management station 11
- Mapping view results 32
- Mappings view, described 26
- Mappings window
  - overview 25
  - views in the 25
- media scan 21
  - changing settings 21
- medical imaging applications 20
- multi-user environments 20
- multimedia applications 20

## N

- Needs Attention status 36
- notes, important 70

## O

- online help systems
  - configuring storage partitions 52
  - configuring storage subsystems 51
  - Enterprise Management window 52
  - event monitor 53
  - miscellaneous system administration 54
  - performance and tuning 55
  - populating a management domain 51
  - protecting data 53
  - recovering from problems 54
  - security 54
  - Subsystem Management window 51, 52
  - using a script editor 52
- Optimal status 36
- overall health status 36

## P

- password protection 22
- performance monitor 47
- physical view, subsystem-management window 17
- populating a management domain 9
  - automatic discovery 9
  - using add device 9
- power outage 21
- protect data in the controller cache memory 20

## Q

- quick reference status
  - Contacting Device 36
  - Fixing 36
  - Optimal 36
  - Unresponsive 36

## R

- RAID level
  - and channel protection 22
  - application behavior 49
  - configurations 19
  - data redundancy 19
  - described 19
- RAID level 0
  - described 20
  - drive failure consequences 20
- RAID level 1
  - described 20
  - drive failure consequences 20
  - redundancy 19
- RAID level 3
  - described 20
  - drive failure consequences 20
  - redundancy 19
- RAID level 5
  - described 20
  - drive failure consequences 20
  - redundancy 19
- README file 1

- reconstruction 21
- Recovery Guru
  - described 41
  - Recovery Procedure 42
  - Summary area 42
- redundancy of Fibre Channel Arbitrated Loops 22
- related publications x

## S

- segment size 50
- sender address configuration 38
- service summary 65
- Simple Network Management Protocol (SNMP)
  - traps 12
- SM7agent disk space requirements 15
- software components
  - RDAC 15
  - SM7agent 15
  - SM7client 14
- start percentage, cache flush 21
- stop percentage, cache flush 21
- storage management software
  - hardware requirements
    - BOOTP server 11
  - installation requirements 11
  - new terminology 2
- storage partition
  - configuration 25
  - creating 29
  - defined
    - mappings 26
    - topology 26
  - described 26
  - major steps to creating 40
  - topology 41
- storage partition topology
  - defined 29
- storage partition topology, defined 26
- storage partitioning
  - feature key 29
  - getting started 28
  - overview 28
- storage subsystem 11
  - configuration 23
  - password protection configuration 22
  - quick reference status icon 36
  - tuning options available 47
- storage subsystems
  - directly managed 5
  - monitoring the health of 35
- storage-partition mapping preference
  - defined 28
- Subsystem Management window 37
  - component of SM7client 14
  - Help 1
  - monitoring storage subsystems with 35

## T

- task reference 51
- topological elements, when to define 26
- Topology view
  - described 26
  - host group with specific mappings 31
  - results 31
- trademarks 70
- tuning storage subsystems 47

## U

- unconfigured and free-capacity nodes 23
- UNIX BOOTP server 11
- Unresponsive status 36

## W

- write cache mirroring
  - described 20
  - how to enable 20
- write caching
  - and data loss 20
  - and performance 20
  - enabling 49









Part Number: 25P1661



Printed in the United States of America  
on recycled paper containing 10%  
recovered post-consumer fiber.

(1P) P/N: 25P1661

