IBM TotalStorage FAStT Storage Manager

Concepts Guide

IBM TotalStorage FAStT Storage Manager
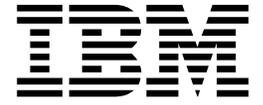
Concepts Guide

> **Note:**
>
> Before using this information and the product it supports, be sure to read the general information under "Notices" on page 115.

First Edition (September 2003)

This edition replaces version 25P1661.

Contents

Figures

Tables

About this document

This document provides the conceptual framework to help you understand the IBM®
TotalStorage® FAStT Storage Manager Version 8.4 for the following
operating-system environments:

- Microsoft® Windows NT®
- Windows® 2000
- Windows® Server 2003
- Novell NetWare
- Sun Solaris
- Hewlett-Packard HP-UX
- IBM AIX®
- Red Hat Linux
- VMWare ESX server

Use this guide to better understand the storage manager software and to perform
the following tasks:

- 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

Who should read this document

This document is intended for system administrators and storage administrators
who are responsible for setting up and maintaining the storage subsystem. Readers
should have knowledge of redundant array of independent disks (RAID), small
computer system interface (SCSI), and fibre-channel technology. They should also
have working knowledge of the applicable operating systems that are used with the
management software.

FAStT installation process overview

Figure 1 on page xii gives an overview of the FAStT hardware and the FAStT
Storage Manager software installation process. Lined arrows in the flow chart
indicate consecutive steps in the hardware and software installation process.
Labeled arrows indicate which current documents provide detailed information about

those steps.

Install Process

- Plan Installation
- Install Storage Server/RAID Controller Enclosure(s) in Rack
- Install Storage Expansion Unit(s)
- Make FC Connections
- Set Link Speed (1 Gb only)
- Determine Management Method
  - Out-of-Band
  - In-Band
- Install Network Hardware; Prepare Network Connection

* For pSeries Server and 6227 or 6228 HBA use only

Documentation

- FC Planning and Integration: User's Guide and Svc Info
- FAStT Storage Manager Concepts Guide

- FAStT Storage Svr Installation Guide
- FAStT RAID Controller Enclosure Unit Install and User's Guide

- FAStT Storage Exp Units Install and User's Guides

- FAStT and HBA Install and User's Guides
- FAStT Fibre Channel Storage Server Installation Guides
- Fibre Channel Cabling Instructions

- Connect Power and Start Server
- Verify Server operation w/ LEDs
- Prepare for Installation of SM Software
- Install and Verify SM SW on Host and Workstation
- Complete SM SW Installation
- Configure Storage Hardware
- Configure Storage Subsystems on Host

- FAStT Storage Server Installation Guides
- Fibre Channel Hardware Maintenance and Problem Determination Guide

- FAStT Storage Manager Installation and Support OS Guides
- FAStT Storage Manager Copy Services User's Guide

- Online Help

Figure 1. Installation process flow by current publications

FAStT documentation

The following table presents an overview of the FAStT Storage Manager product library. The table lists documents included in the library and what common tasks they address. Click on active links in the tables to access those documents currently available on the Internet. You can access documentation for the other FAStT products at the following Web site:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt/

FAStT Storage Manager Version 8.4 library

Table 1 associates each document in the FAStT Storage Manager library with its related common user tasks.

Table 1. TotalStorage FAStT Storage Manager Version 8.4 titles by user tasks

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM TotalStorage FAStT Storage Manager 8.4 Library Guide and Common Index, GC26-7600	✔	✔	✔	✔	✔	✔
IBM TotalStorage FAStT Storage Manager 8.4 Installation and Support Guide for Intel-based Operating System Environments, GC26-7589	✔		✔	✔		
IBM TotalStorage FAStT Storage Manager 8.4 Installation and Support Guide for AIX®, UNIX, and Solaris, GC26-7593	✔		✔	✔		
IBM TotalStorage FAStT Storage Manager Copy Services User's Guide, GC26-7561	✔		✔	✔	✔	
IBM FAStT Storage Manager Script Commands (see product CD)				✔		
IBM TotalStorage FAStT Storage Manager Concepts Guide, GC26-7560	✔	✔	✔	✔	✔	✔

How this document is organized

Chapter 1, "Introduction," on page 1 provides an introduction of the IBM FAStT Storage Manager Version 8.4, compares storage subsystem management methods, and describes the FAStT Storage Manager Enterprise Management and Subsystem Management windows.

Chapter 2, "Protecting your data," on page 39 describes the various data protection features of the FAStT storage servers. These features include input/output (I/O) data path failover support, Media Scan, and copy services.

Chapter 3, "Configuring storage subsystems," on page 57 describes the frequently-used functions of the IBM FAStT Storage Manager. These functions include configuring arrays and logical drives, and mapping these logical drives to separate hosts when the Storage Partitioning premium feature is enabled.

Chapter 4, "Maintaining and monitoring storage subsystems," on page 69, describes how to monitor 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 a failure.

Chapter 5, "Tuning storage subsystems," on page 85 discusses the tuning options that are available in IBM FAStT Storage Manager Version 8.4.

Chapter 6, "Critical event problem solving," on page 89 provides a list of all the critical events that the storage management software sends if a failure occurs. The list includes the critical event number, describes the failure, and refers you to the procedure to correct the failure.

Appendix A, "Online help task reference," on page 103 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.

Appendix B, "Additional FAStT documentation," on page 109 provides a bibliography of additional documentation.

Notices that this document uses

This document contains the following notices designed to highlight key information:
- **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.

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM and 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, what to do if you experience a problem with your IBM @server xSeries™ or IntelliStation® system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you take these steps to try to solve the problem yourself:
- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system.
- Go to the IBM Support Web site at:

  www.ibm.com/pc/support/

to check for technical information, hints, tips, and new device drivers.
* Use an IBM discussion forum on the IBM Web site to ask questions.

You can solve many problems without outside assistance by following the
troubleshooting procedures that IBM provides in the online help or in the
publications that are provided with your system and software. The information that
comes with your system also describes the diagnostic tests that you can perform.
Most xSeries and IntelliStation systems, operating systems, and programs come
with information that contains troubleshooting procedures and explanations of error
messages and error codes. If you suspect a software problem, see the information
for the operating system or program.

Using the documentation

Information about your xSeries or IntelliStation system and preinstalled software, if
any, is available in the documentation that comes with your system. That
documentation includes printed documents, online documents, readme files, and
help files. See the troubleshooting information in your system documentation for
instructions for 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, follow the instructions on the Web site at:

www.ibm.com/pc/support/

Also, to order publications through the IBM® Publications Ordering System, go to
the Web site at:

www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi/

Web sites

On the World Wide Web, the IBM Web site has up-to-date information about IBM
xSeries and IntelliStation products, services, and support. For IBM xSeries
information, go to the Web site at:

www.ibm.com/eserver/xseries/

For IBM IntelliStation information, go to the Web site at:

www.ibm.com/pc/intellistation/

For service information for your IBM products, including supported options, go to
the Web site at:

www.ibm.com/pc/support/

If you click **Profile** from the support page, you can create a customized support
page. The support page has many sources of information and ways for you to solve
problems, including how to:
* Diagnose problems, using the IBM Online Assistant
* Download the latest device drivers and updates for your products
* View Frequently Asked Questions (FAQ)
* View hints and tips to help you solve problems

- Participate in IBM discussion forums
- Set up e-mail notification of technical updates about your products

For the latest information about operating system and HBA support, clustering support, storage area network (SAN) fabric support, and Storage Manager feature support, see the IBM TotalStorage FAStT Interoperability Matrix at the following Web site:

www.storage.ibm.com/disk/fastt/supserver.htm

Software service and support

Through the IBM Support Line, for a fee, you can get telephone assistance with usage, configuration, and software problems with xSeries servers, IntelliStation workstations and appliances. For information about which products are supported by Support Line in your country or region, go to the Web site at:

www.ibm.com/services/sl/products/

For more information about Support Line and other IBM services, go to the Web site at:

www.ibm.com/services/

or, for support telephone numbers, go to the Web site at:

www.ibm.com/planetwide/

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. For support telephone numbers, go to the Web site at:

www.ibm.com/planetwide/

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

How to send your comments

Your feedback is important to help us provide the highest quality of information. If you have any comments about this document, you can submit them in one of the following ways:

- E-mail

  Submit your comments electronically to:

  starpubs@us.ibm.com

  Be sure to include the name and order number of the document and, if applicable, the specific location of the text that you are commenting on, such as a page number or table number.

- Mail or fax

  Fill out the Readers' Comments form (RCF) at the back of this document and return it by mail or fax (1-800-426-6209) or give it to an IBM representative. If the RCF has been removed, you may address your comments to:

International Business Machines Corporation
RCF Processing Department
Dept. M86/Bldg. 050-3
5600 Cottle Road
San Jose, CA 95193-0001
U.S.A.

Chapter 1. Introduction

This concepts guide provides the conceptual framework that is necessary to understand the terminology and features of the IBM FAStT Storage Manager Version 8.4. Before you install 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, for the latest installation and user information about the storage management software and hardware components, go to the following Web site:

  www.storage.ibm.com/support/

  From this Web site, choose FAStT Storage Servers.

- 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 8.4.

When you finish the entire installation process, refer to the following online help systems, which contain information that is 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 8.4, click **Help** on the toolbar, or press F1.

Machine types and supported software

Table 2 provides a list of machine types and supported storage management software.

Note: Product firmware must be at version 4.x to be compatible with Storage Manager Version 8.40.

Table 2. Machine types and supported storage management software

Product name	Machine type	Model	Firmware version	Supported storage management software version
Netfinity Fibre Channel RAID Controller Unit	3526	1RU 1RX	4.x	7.0, 7.01, 7.02, 7.10, 8.0, 8.2, 8.21, 8.3, 8.4
IBM FAStT200 Storage Server	3542	1RU 1RX	4.x, 5.20.xx.xx, 5.30.xx.xx	7.02, 7.10, 8.0, 8.2, 8.21, 8.3, 8.4
IBM FAStT200 High Availability Storage Server	3542	2RU 2RX	4.x, 5.20.xx.xx, 5.30.xx.xx	7.02, 7.10, 8.0, 8.2, 8.21, 8.3, 8.4

Table 2. Machine types and supported storage management software (continued)

Product name	Machine type	Model	Firmware version	Supported storage management software version
Netfinity FAStT500 RAID Controller Enclosure Unit	3552	1RU 1RX	4.x, 5.00.xx.xx 5.20.xx.xx 5.21.xx.xx, 5.30.xx.xx	7.0, 7.01, 7.02, 7.10, 8.0, 8.2, 8.21, 8.3, 8.4
IBM TotalStorage FAStT700 Fibre Channel Storage Server	1742	1RU 1RX	5.00.xx.xx 5.20.xx.xx 5.21.xx.xx, 5.30.xx.xx 5.4x.xx.xx	8.0, 8.2, 8.21, 8.3, 8.4
IBM TotalStorage FAStT900 Fibre Channel Storage Server	1742	90X, 90U	5.30.xx.xx 5.4x.xx.xx	8.3, 8.4
IBM TotalStorage FAStT600 Fibre Channel Storage Server	1722	60X, 60U	5.33.xx.xx 5.4x.xx.xx	8.3, 8.4
IBM TotalStorage FAStT600 Turbo Fibre Channel Storage Server	1722	60X, 60U	5.33.xx.xx 5.4x.xx.xx	8.4

Note: Controller firmware versions regulate the function of any given Storage Manager client code version release. You might need to purchase the appropriate firmware upgrade option to enable the functions of your particular Storage Manager client code version release. Storage subsystem firmware must be at version 04.00.02.xx or higher to be managed with Storage Manager 8.0, 8.2, 8.21, 8.3, or 8.4.

Terms to know

If you are upgrading from a previous version of the storage management software, you will find that some of the terms that you are familiar with have changed. It is important that you familiarize yourself with the new terminology. Table 3 provides a list of some of the old and new terms.

Table 3. Old and new terminology

Term used in previous versions	New term
RAID module	Storage subsystem
Drive group	Array
Logical unit number (LUN) (See note)	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

Table 3. Old and new terminology (continued)

Term used in previous versions	New term

Note: 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 one of 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 that is installed on a host computer that is directly attached to the storage subsystem through the fibre-channel 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 tasks:
>
>- 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 document.
2. A host computer can also function as a management station.

New features and enhancements

The following features are new to the IBM FAStT Storage Manager Version 8.4:

- Support for IBM TotalStorage FAStT600 Turbo

  FAStT600 Turbo consists of controllers with 1 GB cache each, Storage Manager version 8.4, eight storage partitions, and up to seven EXP700 drive expansion attachments.

- VolumeCopy

  VolumeCopy allows you to create a complete copy of a logical drive within the same storage array. The feature is designed as a system management tool for tasks such as copying data to arrays on other drives for hardware upgrades or performance management, data backup, and restoring FlashCopy logical drive data within the same storage array. This feature is supported on FAStT600, FAStT600 Turbo, FAStT700 and FAStT900.

In addition to these new features, the following enhancements have been made to the software:

- 256 logical drives per storage mapping partition

  Important: Windows NT 4.0, Sun Solaris with RDAC, Novell NetWare 5.1, and HP-UX 11.0 are restricted to 32 volumes per partition. There are limitations on other operating systems as well. Refer to the following table for more information:

Table 4. Number of LUNs supported per operating system

Operating environment	Highest number of LUNs the operating system will look for	Highest number of LUNs supported per storage mapping partition
Windows NT	254	32
Windows 2000	254	64 (See note 5)
Windows Server 2003	254	255
Sun Solaris	255 (See note 1)	256
HP-UX 11.0	32 (with VSA)	32
HP-UX 11.11	127	128
IRIX ~1000 (See note 2)	256	256
AIX 31 /1000+ (See note 3)	256	256
Novell Netware 5.1	31	32
Novell Netware 6.0	255	256
Linux (Red Hat)	255	128 (See note 4)
Linux (SuSE)	255	256
VMWare ESX 2.0	127	128

Notes:

1. The number of logical drives that is supported is driver-dependent.
2. This is an operating system installation option. The default might be less.
3. The limit is 32 through the AIX classic RDAC driver but the operating system limit is higher.
4. The limit is 128 devices due to an operating system limitation.
5. The highest number of LUNs supported per storage mapping partition on Windows 2000 is restricted due to registry limitations and the possibility of file system corruption.

- Persistent reservations

  Persistent reservations can be used to preserve logical drive reservations and registrations and to prevent other hosts from accessing the logical drive. Persistent reservations perform the following tasks:

  – Reserve access across multiple host ports
  – Provide various levels of access control
  – Offer the ability to query the storage array about registered ports and reservations
  – Provide for persistence of reservations in the event of a storage system power loss

- Veritas Volume Manager with Dynamic Multi-Pathing (DMP)

  Support for the Veritas Volume Manager with Dynamic Multi-Pathing (DMP) on Solaris has been enhanced. A new host type, Solaris with Veritas DMP, has been added to give more flexibility when you configure the host-based multi-path drivers. After you have installed it, the new Array Support Library (ASL) will provide information to the Veritas Volume Manager to set up path associations for the failover driver. For more information on Veritas Volume Manager with Dynamic Multi-Pathing (DMP), see the *IBM TotalStorage FAStT Storage Manager 8.4 Installation and Support Guide for AIX®, UNIX, and Solaris*.

- Password failure reporting and lockout

The password failure reporting and lockout feature provides additional security for storage servers with a configured password. Any failed attempt to access the storage array will be reported as an informational MEL event. If the password is incorrectly entered 10 times within 10 minutes, both controllers will enter lockout mode, preventing unauthorized access to your storage array.

- The FlashCopy feature is supported on the FAStT600 and FAStT600 Turbo storage servers
- User control of network parameters

  You can now change the network parameters, for example, you can change the IP address and subnet mask for the controller, and enable or disable the Rlogin capability and BOOTP.
- Support for 2 TB arrays or larger
- Support for the creation of logical drives with a segment size of 512K using the script engines or SMcli
- New script engines, SMcli commands and enhancements of existing commands

  These changes improve script engine error reporting.
- Support on each storage subsystem, for a human-readable, append-only log that updates automatically

  This log file that can be used by IBM support personnel to recover from storage subsystem catastrophic failures.

Notes:

- The controller firmware must be at level 05.4x..xx.xx or higher and the Storage Manager client software version must be 8.4x in order to use these new features and enhancements.
- These new features and enhancements are not available for FAStT200 and FAStT500 because the highest supported firmware version for these storage subsystems is version 05.30.xx.xx.

The following premium features can be enabled by purchasing a premium feature key:

- The Storage Partitions feature groups logical drives into sets, called *storage partitions*, to consolidate storage and reduce storage management costs. The Storage Partitioning feature is enabled by default for some FAStT storage server models and configurations. For more information about storage partitions, see "Storage partitioning" on page 60.
- FlashCopy® supports creating and managing FlashCopy logical drives. A FlashCopy logical drive is a logical point-in-time image of another logical drive, called a *base logical drive*, that is in the storage subsystem. A FlashCopy is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space. For more information about FlashCopy, see "FlashCopy" on page 50 or see the *IBM TotalStorage FAStT Storage Manager 8.4 Copy Services Guide*.
- The VolumeCopy feature is a firmware-based mechanism that is used to copy data from one logical drive (the *source* logical drive) to another logical drive (the *target* logical drive) in a single storage subsystem. The VolumeCopy feature can be used to copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives, to back up data, or to restore FlashCopy logical drive data. The VolumeCopy feature includes a Create Copy Wizard that is used to assist in creating a VolumeCopy, and a Copy Manager that is used to monitor VolumeCopies after they have been created. For more information about

VolumeCopy, see "VolumeCopy" on page 51 or see the *IBM TotalStorage FAStT Storage Manager 8.4 Copy Services Guide*

- The Remote Mirror Option provides real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Mirror Option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations. For more information about the Remote Mirror Option, see "Remote Mirror option" on page 52 or see the *IBM TotalStorage FAStT Storage Manager 8.4 Copy Services Guide*.

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 might span multiple physical enclosures depending on the number of drives and the RAID controller technology. The FAStT500, FAStT700, and FAStT900 server configurations consist of an independent RAID controller enclosure and at least one drive expansion enclosure (EXP500 and EXP700).The FAStT200, FAStT600 and FAStT600 Turbo servers integrate the drive expansion enclosure and the RAID controller function in the same physical enclosure. The FAStT200, FAStT600 and FAStT600 Turbo can use external drive enclosures. Figure 2 shows an example of the minimum FAStT500 storage subsystem physical components.
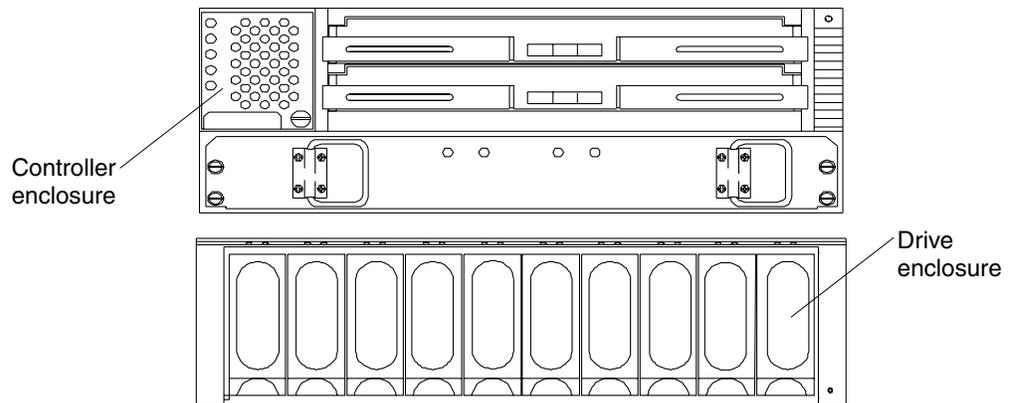
Figure 2. Storage subsystem physical components

The maximum number of drives and drive expansion enclosures that a RAID controller can support depends on the model of the RAID storage servers. See the User's Guides for your storage server for the maximum number of drives and drive expansion enclosures that are supported per storage server.

Table 5 describes the storage subsystems physical components.

Table 5. Storage subsystem physical components

Component	Description
Drive	An electromagnetic mechanical device that provides the physical data storage media.
Drive expansion enclosure	An enclosure that contains drives, power supplies, fans, environmental service monitors, assembly, and other supporting components.

Table 5. Storage subsystem physical components (continued)

Component	Description
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.

The physical disk capacity of the storage server is divided into arrays and logical drives. These are recognized by the operating system as unformatted physically attached disks. Each logical component can be configured to meet data availability and I/O performance needs. Table 6 describes the storage subsystem logical components.

Table 6. Storage subsystem logical components

Component	Description
Array	An array is 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. The number of drives that can be grouped together into an array depends on the hard drive capacity and the FAStT storage server firmware version. Each array can be divided into 1 - 256 logical drives. For more information about arrays, see the FAStT Storage Manager Installation and Support Guide for your operating system.
Logical drive	A logical drive is a logical structure that you create to store data in the FAStT storage server. A logical drive is a contiguous subsection of an array that is configured with a RAID level to meet application needs for data availability and I/O performance. The operating system sees the logical drive as an unformatted drive.
Storage partition	A storage partiton is a logical identity that consists of one or more storage subsystem logical drives. The storage partition is shared with host computers that are part of a host group or is accessed by a single host computer. Use the Storage Partition premium feature key to enable a storage partition. The number of storage partitions that can be defined depends on the model of the FAStT storage servers. See Table 7 on page 8 to determine whether the Storage Partition feature is enabled by default and the maximum number of storage partitions that can be defined for a given FAStT storage server.
Free capacity	Free capacity is a contiguous region of unassigned capacity on a designated array. You will use free capacity to create one or more logical drives. **Note:** In the Subsystem Management window Logical view, free capacity is displayed as free capacity nodes. Multiple free capacity nodes can exist on an array.

Table 6. Storage subsystem logical components (continued)

Component	Description
Unconfigured capacity	Unconfigured capacity is the capacity in the storage subsystem from drives that are not assigned to any array. You will use this space to create new arrays. **Note:** In the Subsystem Management window Logical view, unconfigured capacity is displayed as an unconfigured capacity node.

Table 7 indicates whether the Storage Partition feature is enabled by default and the maximum number of storage partitions that can be defined for a given FAStT storage server.

Table 7. Storage partitioning specifications per FAStT server model

FAStT server models	Storage Partitioning enabled by default	Maximum number of defined storage partitions	Available storage partition purchase options
FAStT200 (3542 1RU/1RX)	Yes	16	None
FAStT600 (1722 60U/60X)	No	16	4, 8, 4 - 8, 16
FAStT600 Turbo (1722 60U/60X)	Yes (8 partitons)	64	8 - 16, 8 - 64, 16 - 64
FAStT500 (3552 1RU/1RX)	Yes	64	None
FAStT700 (1742 1RU/1RX)	Yes	64	None
FAStT900 (1742 90U/90X)	Yes (16 partitions. See note.)	64	16 - 64

Note: The FAStT900 ships with a default 16 partitions enabled. You can buy a 16 - 64 partition upgrade.

Software components

Table 8 describes the IBM TotalStorage FAStT Storage Manager Version 8.4 software components.

Table 8. Storage Manager software components

Component	Description
Storage Manager 8.4 client (SMclient)	The Storage Manager 8.4 client (SMclient) component provides the graphical user interface (GUI) for managing storage subsystems through the Ethernet network or from the host computer. The SMclient contains two main components: • **Enterprise Management**. You can use the Enterprise Management component to add, remove, and monitor storage subsystems within the management domain. • **Subsystem Management**. You can use the Subsystem Management component to manage the components of an individual storage subsystem. The Storage Manager 8 client is called *thin* because it only provides an interface for storage management based on information that is supplied by the storage subsystem controllers. When you install the SMclient 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 validates and runs the commands and provides the status and configuration information that is sent back to the SMclient.

Table 8. Storage Manager software components (continued)

Component	Description
Storage Manager 8.4 host agent (SMagent)	The Storage Manager 8.4 agent (SMagent) 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. 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. 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 the requests to the controllers in the storage subsystem through the fibre-channel I/O path. **Note:** Host computers that have the host agent software installed are automatically discovered by the storage management software. They are displayed in the device tree in the Enterprise Management window along with their attached storage subsystems. A storage subsystem might be duplicated in the device tree if you are managing it through its Ethernet connections and it is attached to a host 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.

Table 8. Storage Manager software components (continued)

Component	Description
Redundant disk array controller (RDAC) multipath driver	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, such as a cable or the controller itself, fails, the RDAC multipath driver transfers ownership of the logical drives that are assigned to that controller to the other controller in the pair. **Note:** The RDAC driver is not available for Novell Netware and Linux operating systems. In the Novell Netware environment, an IBM Storage Area Network (IBMSAN) driver is used in place of RDAC. In the Linux environment, a failover capable host bus adapter (HBA) driver with a FAStT-MSJ program and Qlogic QLremote agent is used to provide multipath failover function and to set up good and failover I/O paths. Some operating systems have built-in fibre-channel I/O path failover drivers and do not require this multipath driver.
The IBM Storage Area Network (IBMSAN) driver for storage manager version 8.3 and earlier (controller firmware 05.3x.xx.xx and earlier)	For Storage Manager version 8.3 and earlier in a Novell NetWare environment, use the IBM Storage Area Network (IBMSAN) driver 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 HBA driver. In case of a failure along the I/O path, the driver sends the I/O requests through the alternate path and auto-logical drive transfer (ADT) moves the logical drive within the FAStT storage server or controller enclosure. The controller ADT function must be enabled for multipath I/O support. To enable the ADT function, set the correct HBA ports that are installed in the Novell Netware servers. **Note:** For Storage Manager version 8.4 and later, Novell Netware provides native multipath support. The IBMSAN driver should not be installed if you are going to use the native Novell Netware multipath driver. ADT must be enabled when you use the native Novell Netware multipath driver for LUNs failover and failback.
Storage Manager utility (SMutil)	Use the Storage Manager utility package to register and map new logical drives to the operating system. SMutil is installed on all host computers. The host computers are attached to the storage subsystem through the fibre-channel connection. The SMutil package contains the following two components: • **Hot Add utility:** The Hot Add utility enables you to register newly-created logical drives with the operating system. • **SMdevices utility:** You can use the SMdevices utility to associate storage subsystem logical drives with operating-system device names. **Note:** See the storage management software readme file for all supported operating systems.

Table 8. Storage Manager software components (continued)

Component	Description
Microsoft Virtual Machine (MSVM)	For Storage Manager version 8.3 and earlier, Microsoft Virtual Machine, which is available on the Storage Manager software CD, is required to support the Event Monitor option of the SMclient package and the SMagent package on a host computer running Windows NT 4.0. You must manually install Microsoft Virtual Machine on your management station, host computer, or server as applicable to your configuration.
Linux failover driver, FAStT-MSJ, and QLremote agent	The Linux failover driver is a failover version of the fibre-channel HBA driver. Use the FAStT-MSJ and the QLremote agent to configure the device driver and the host adapters for multipath failover. In case of a failure along an I/O path, the multipath failover driver sends the I/O requests through the alternate path and ADT moves the logical drive within the FAStT storage server from the controller that was in the failed path to the controller that is in the good path. **Note:** The controller ADT function must be enabled for multipath I/O support. To enable the ADT function, set the correct HBA ports that are installed in the Linux servers.

System requirements

This section provides detailed information about the hardware, software, and storage management architecture for IBM FAStT Storage Manager Version 8.4.

Hardware requirements

Table 9 lists the hardware that is required to install Storage Manager 8.4.

Table 9. Storage management architecture hardware components

Component	Description
Management station (one or more)	A management station is a computer that is connected through an Ethernet cable to the host computer or directly to the storage controller. • Monitor setting of 1024 x 768 pixels with 64 000 colors. The minimum display setting that is allowed is 800 x 600 pixels with 256 colors. • Hardware-based Windows acceleration. Desktop computers that use system memory for video memory are not preferred for use with the storage management software. **Important:** Many PC-based servers are not designed to run graphic-intensive software. If your server has difficulty running the storage management software smoothly without video artifacts, you might need to upgrade the server video adapter.

Table 9. Storage management architecture hardware components (continued)

Component	Description
Network-management station (optional, for SNMP traps)	A network-management station is a computer with installed SNMP-compliant network-management software. It 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 DHCP/BOOTP or BOOTP-compatible server is a server that is used to assign the network-specific information such as an internet protocol (IP) address and host computer name for each controller. **Note:** You do not need to set up the DHCP/BOOTP server if the static IP addresses or default IP addresses of the controllers are used, or if you are managing all storage subsystems through the fibre-channel I/O path using a host agent.
Host bus adapters (HBAs)	Host bus adapters are one or more adapters that are installed in the host server that provide the fibre-channel interface port or ports for the fibre connection between the storage subsystem and the host server.
Fibre-channel switches	Fibre-channel switches are used if there are more host servers that need to access the storage subsystem than the available number of physical fibre-channel ports on the storage subsystem. **Note:** For FAStT500, FAStT700, and FAStT900 storage servers, if one of the two ports of a host minihub is connected to the fibre-channel switch, the other minihub port must be left unconnected (open).
Host computer	A host computer is a computer that is running one or more applications that accesses the storage subsystem through the fibre-channel I/O data connection.
Storage subsystem and storage server (one or more)	The storage subsystem and storage server are storage entities that are 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. Management stations on the network can then remotely access the storage management software.

Firmware requirements

To use all of the features of Storage Manager 8.4, your storage subsystem must contain controllers with firmware version 5.4x.xx.xx for machine types 1722 and 1742. To manage storage subsystems with Storage Manager 8.4, your storage subsystems must contain controllers with firmware versions 04.00.02.xx or higher.

The highest supported firmware for machine types 3542 and 3552 is version 05.3x.xx.xx.

Important: You must purchase the option to upgrade your storage subsystem controller firmware from version 05.3x.xx.xx or earlier to 05.4x.xx.xx. Contact your IBM technical-support representative for more information.

Storage Manager 8.4 also supports four premium features: Storage Partitioning, VolumeCopy, FlashCopy and the Remote Mirror Option. For more information about these premium features, contact your IBM reseller or IBM marketing representative. For more information about FlashCopy, VolumeCopy or the Remote Mirror Option, see the *IBM FAStT TotalStorage Copy Services User's Guide*. The Storage Partitioning feature is enabled by default for some FAStT storage server models and configurations.

Storage subsystem management

IBM FAStT Storage Manager provides two methods for managing storage subsystems: the host-agent (in-band) management method and the direct (out-of-band) management method. A storage subsystem receives data, or I/O, from the application host server over the fibre-channel I/O path. The FAStT storage management software can be installed in the host server or in a management workstation that is in the same network as the host computer (in-band management method) or the storage subsystem (out-of-band management method). Depending on your specific storage subsystem configurations, you can use either or both methods.

Storage subsystem management include the following activities:
- 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 the storage subsystem for optimal application performance

Direct (out-of-band) management method

When you use the direct (out-of-band) management method, you manage storage subsystems directly over the network through the Ethernet connection to each controller. To manage the storage subsystem through the Ethernet connections, you must define the 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 3 on page 16.

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 SMclient to manage storage subsystems that are connected to a host computer running an operating system that is supported by IBM FAStT Storage Manager Version 8.4.
- 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.
- You can manage and troubleshoot the storage subsystem when there are problems with the fibre-channel links.

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, see the installation and support guide for your operating system.
- You can avoid DHCP/BOOTP server and network tasks by assigning static IP addresses to the controller or by using a default IP address.

To assign a static IP address, see Retain Tip H171389 "Unable To Setup Networking Without DHCP/BOOTP" at the following Web site:

www.ibm.com/support/

Table 10 lists the default settings for storage subsystem controllers that have firmware version 05.00.xx. or later:

Table 10. Default settings for controllers with firmware version 05.00.xx or later

Controller	IP address	Subnet mask
A	192.168.128.101	255.255.255.0
B	192.168.128.102	255.255.255.0

Figure 3 shows a system in which the storage subsystems are managed through the direct (out-of-band) management method.

Storage subsystems

Controller

Controller

Host computer

Fibre-channel
I/O path

Controller

Controller

Storage subsystems

Ethernet

Network-management connection

Management station
(one or more)
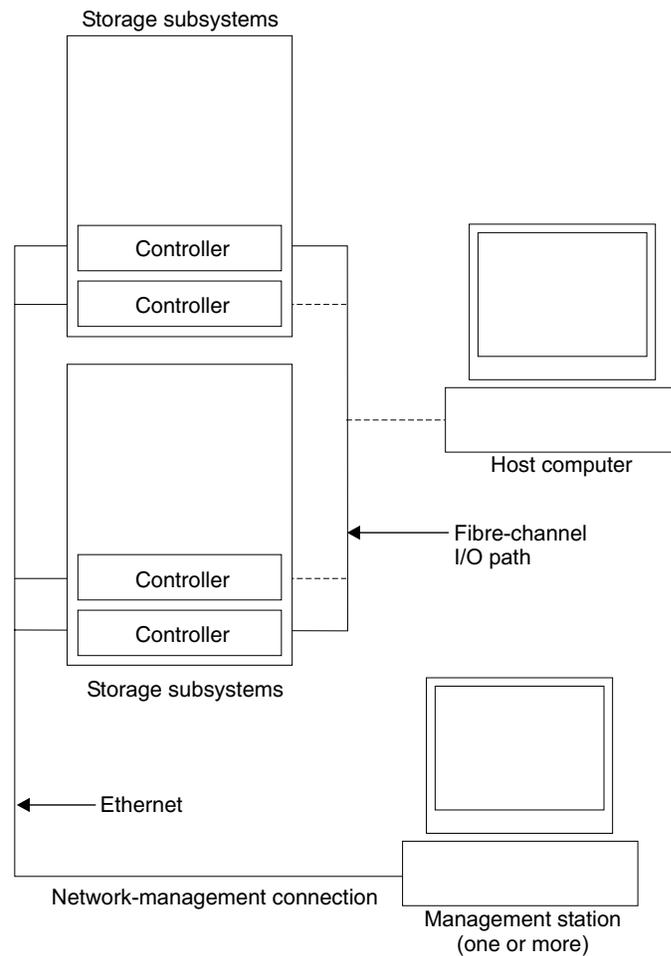
SJ001087

Figure 3. Direct (out-of-band) managed storage subsystems

Host-agent (in-band) management method

When you use the host-agent (in-band) management method, the controllers in the storage subsystem are managed through the host-agent 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. The management software can be installed in the host or in the management station that is connected to the host through the Ethernet network connection. 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:

- The host agent requires a special logical drive, called an *access volume*, to communicate with the controllers in the storage subsystem. Therefore, you are limited to configuring one less logical drive than the maximum number that is allowed by the operating system and the host adapter that you are using.

  Important: If your host already has the maximum number of logical drives configured, either use the direct management method or give up a logical drive for use as the access logical drive.

  Note: The access logical drive is also referred to as the *Universal Xport Device*.
- If the connection through the fibre channel is lost between the host and the subsystem, the subsystem cannot be managed or monitored.

Figure 4 on page 18 shows a system in which the storage subsystems are managed through the host agent (in-band) management method.

Storage subsystems

Controller

Controller

Running the
host-agent software

Host computer

Fibre-channel
I/O path

Controller

Controller

Storage subsystems

Network

Note: The host can act as a
management station also.

Management station
(one or more)
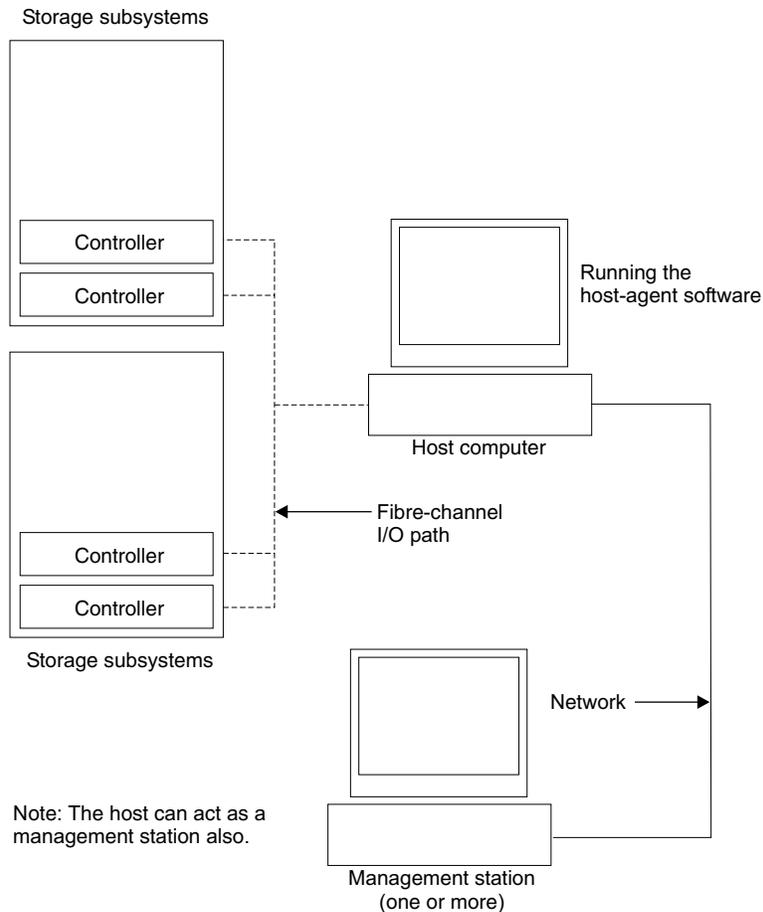
SJ001089

Figure 4. Host-agent (in-band) managed storage subsystems

Reviewing a sample network

Network A in Figure 5 on page 19 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 SNMP traps
- Host computer that is connected to a storage subsystem through a fibre-channel
  I/O path
- Management station connected by an Ethernet cable to the storage subsystem
  controllers

Network B in Figure 5 on page 19 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 an Ethernet cable to the host
  computer
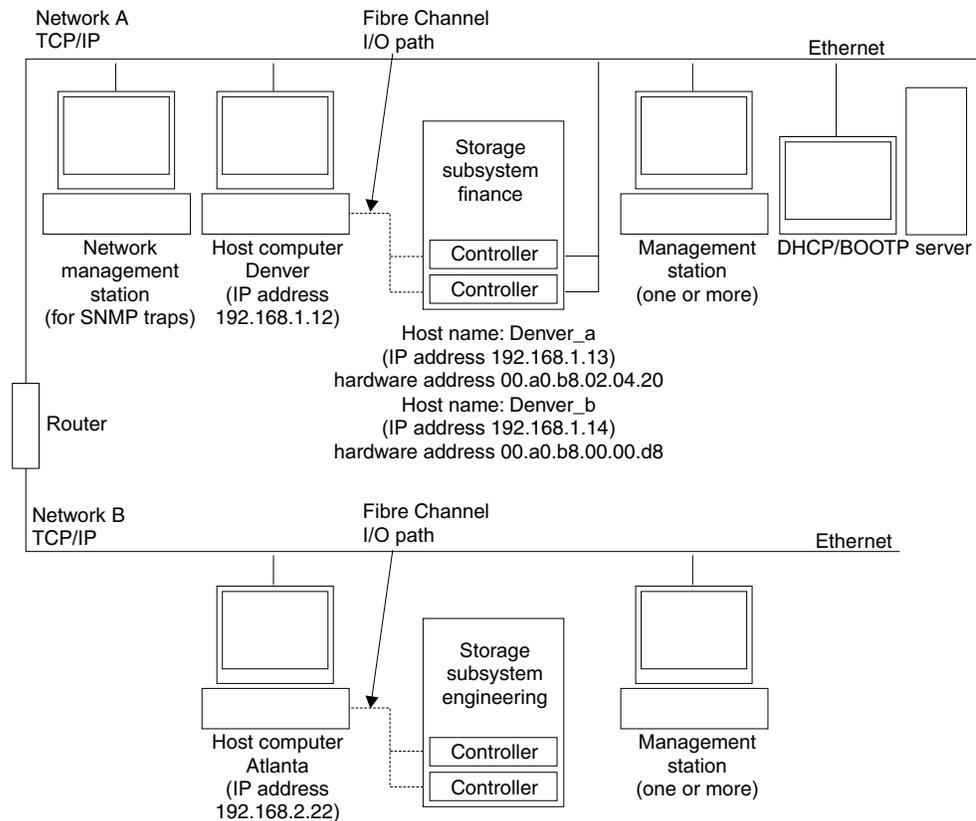
Figure 5. Sample network using direct and host-agent managed storage subsystems

Managing coexisting storage subsystems

Storage subsystems are *coexisting storage subsystems* when the following conditions are met:

- Multiple storage subsystems with controllers are running different versions of the firmware.
- These storage subsystems are attached to the same host.

For example, a coexisting situation exists when you have a new storage subsystem with controllers that are running firmware version 05.4x.xx.xx, and the storage subsystem is attached to the same host as one or more of the following configurations:

- A storage subsystem with controllers running firmware versions 04.00.xx.xx through 04.00.01.xx, which is managed by a separate management station with Storage Manager Version 7.10
- A storage subsystem with controllers running firmware versions 04.00.02.xx through 05..3x.xx.xx, which is managed with Storage Manager Version 8.4

Important: The common host must have the latest level (version 8.4) of RDAC and SMagent installed. The latest 05.4x.xx.xx firmware is available free of charge for download from the IBM support Web site along with all the fixes and software patches. You must upgrade all storage subsystem 05.2x.xx.xx-level (or earlier) firmware to the current code level.

Figure 6 shows an example of a Windows environment that includes coexisting storage subsystems.

Version 8.4 of the storage-management software must be installed on either a management station or host.

Important: In a co-existing situation, the common host must have Storage Manager Version 8.4 RDAC and SMagent installed.

Management station

Firmware version 4.x (minimum)

Controller

Controller

Host-agent

Host computer

Firmware version 5.x

Controller

Controller

Fibre Channel I/O path

Ethernet

Storage subsystems

Client software

Version 7.10 of the storage-management software must be installed on a management station.*

*Only to manage existing storage subsystems with firmware version 4.00.xx.xx to 4.00.01.xx.

Management station

SJ001058

Figure 6. Sample coexisting storage subsystems

Managing the storage subsystem using the graphical user interface

This section includes information about managing the storage subsystem using the SMclient graphical user interface (GUI) and covers the following topics:

- The Enterprise Management window
- The Subsystem Management window
- Populating the management domain
- The script editor

The two main windows are the Enterprise Management window and the Subsystem Management window. The Enterprise Management window is shown in Figure 7 on page 21. The Subsystem Management window is shown in Figure 9 on page 25.

Enterprise Management window

The Enterprise Management window, as shown in Figure 7 is the first window that opens when you start the storage management software. Use the Enterprise Management window to perform the following tasks:

- 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
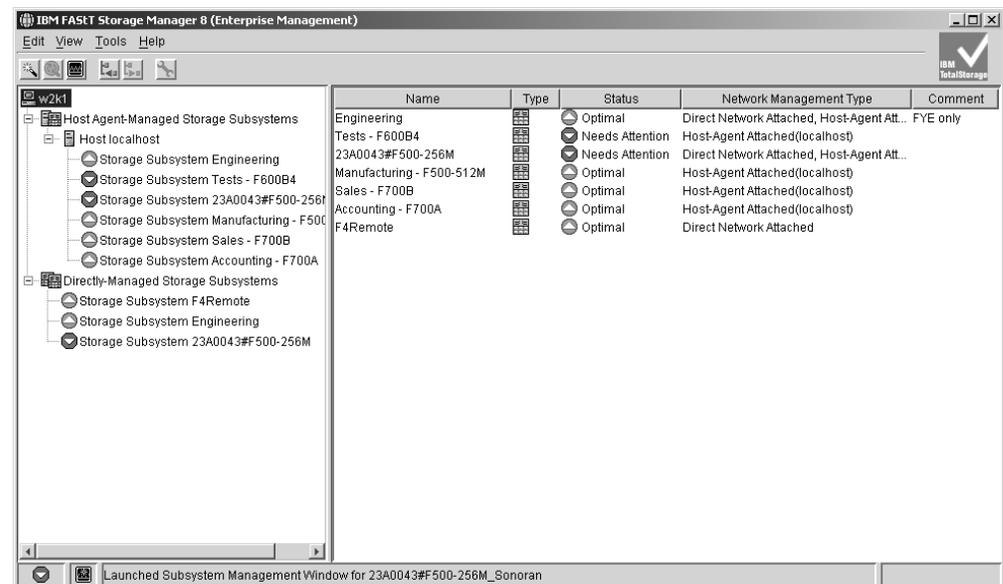
SJ001136

Figure 7. The Enterprise Management window

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. You can also 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 name of the default directory depends on your operating system and firmware version.

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

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 8.

Management station — HFordPC
Host-agent managed — Host Kansas City, Storage Subsystem Engineering
Direct managed — Storage Subsystem Finance
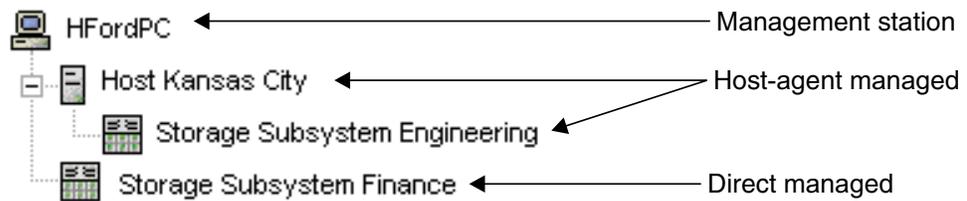
Figure 8. Device tree with a management domain

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 display 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.

Read the following points before you populate a management domain:

- Be sure to specify the IP addresses for both controllers when you add new storage subsystems to existing storage subsystems that are managed using the direct-management method.

- If a given FAStT storage subsystem is listed in the device tree as being both directly-managed and host-agent managed, the FAStT storage manager program will select the directly-managed route to manage the storage subsystem.

  Note: If the FAStT subsystem is seen by the SMclient through both in-band and out-of-band management methods, the subsystem will be displayed in two places in the device tree.

- When you add new storage subsystems to the existing storage subsystems in a SAN that are managed through the host-agent software, you must stop and restart the host-agent service. When the host-agent service restarts, the new storage subsystem is detected. Then, go to the Enterprise Management window and click **Tools → Rescan** to add the new storage subsystems to the management domain.

- If you have a large network, the Automatic Recovery option might take a while to complete. You might also get duplicate storage subsystem entries listed in the device tree if there are multiple hosts in the same network that has a host-agent connection to the storage subsystems. You can remove a duplicate storage management icon from the device tree by using the Remove Device option in the Enterprise Management window.

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

Subsystem Management window

The Subsystem Management window is used to configure and maintain the logical and physical components of a storage array and to view and define volume-to-LUN mappings.

The Subsystem Management window is specific to an individual storage array; therefore, you can manage only a single storage array within a Subsystem Management window. However, you can start other Subsystem Management windows from the Enterprise Management window to simultaneously manage multiple storage arrays.

Note: The features of a particular release of firmware will be accessible when a Subsystem Management window is launched from the Enterprise Management window to manage a storage array. For example, you manage two storage arrays using this software; one storage array has firmware version 5.40 and the other has firmware version 4.x. When you open a Subsystem Management window for a particular storage array, the correct Subsystem Management window version is used. The storage array with firmware version 5.40 will use version 8.40 of the storage management software, and the storage array with firmware version 4.x will use version 7.x. You can verify the version you are currently using by clicking **Help ➔ About in the Subsystem Management Window**.

This bundling of previous versions of the Subsystem Management window provides the flexibility of upgrading the firmware only on selected storage arrays instead of having to perform an upgrade on all storage arrays at once.

Subsystem Management window tabs

The Subsystem Management window consists of two views: the Logical view and the Physical view. These two views are side by side in a single window. Only one of the views can be displayed at a given time in a single Subsystem Management window. The view is selected by clicking on appropriate folder tab. The Subsystem Management window has a second window, the Mappings window.

Table 11. Subsystem Management window tabs

Tabs	Description
Logical/Physical View	The Subsystem Management Window Logical/Physical View contains two panes: the Logical View and the Physical View.
	The Logical View (left pane of Figure 9 on page 25) provides a tree-structured view of logical nodes. This view shows the organization of storage array capacity into arrays and logical drives.
	The Physical View (right pane of Figure 9 on page 25) provides a view of the physical devices in a storage array, such as controller tray and drive tray components.
	Selecting a logical drive or other entity in the Logical View shows you the associated physical components in the Physical View.
	There is a **Components** button in every controller tray and drive tray that, when clicked, presents the status of each component and shows the temperature status.
Mappings View	The Mappings view of the Subsystem Management window contains two views, Topology and Mappings.
	The Topology View provides a tree-structured view of logical nodes related to storage partitions.
	The Mappings View displays the mappings associated with the selected node in the Topology View.
	For more information about the Mappings View, see "Mappings window overview" on page 64.
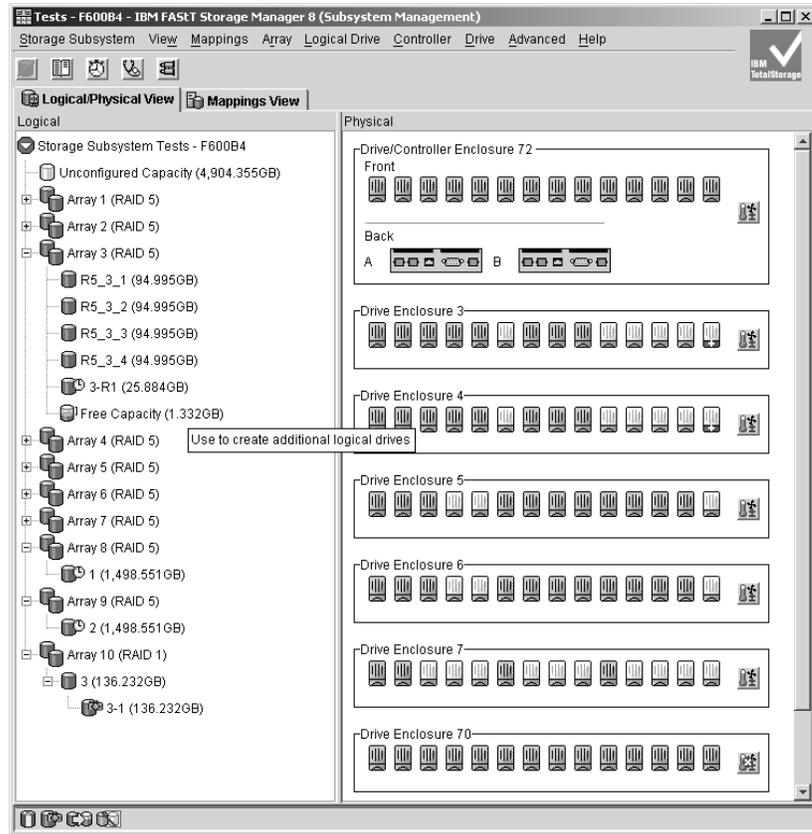
Figure 9. Subsystem Management window Logical View and Physical View

The Subsystem Management window menus

The Subsystem Management window menus are described in Table 12 on page 26. The menus are used to perform storage management operations for a selected storage array or for selected components within a storage array.

Table 12. The Subsystem Management window menus

Menu	Definition
Storage Subsystem	The **Storage Subsystem** menu contains options to perform the following storage array management operations: • Locating functions (locating the storage array or a specific drive channel by flashing indicator lights) • Configuring the storage subsystem • Enabling premium features • Starting the Recovery Guru • Monitoring performance • Downloading firmware and NVSRAM files • Changing various Storage Subsystem settings • Setting controller clocks • Redistributing logical drives • Running Read Link Status diagnostics • Activating or deactivating the Remote Mirror Option • Renaming storage subsystem
View	The **View** menu allows you to perform the following tasks: • Change the display from the Logical/Physical view to the Mappings View • View associated components for a selected node • Find a particular node in the Logical View or Topology View • Locate an associated node in the tree • Access the Event Log or the Storage Subsystem Profile
Mappings	The **Mappings** menu permits you to make changes to or retrieve details about mappings associated with a selected node. The options are: • Define • Change • Move • Replace • Host Port • Show All Host Port Information • Remove • Rename **Note:** You must be in the Mappings View to access the options available in this menu.

Table 12. The Subsystem Management window menus (continued)

Menu	Definition
Array	The **Array** menu presents options to perform the following storage management operations on arrays: • Creating or locating logical drives, changing RAID level or controller ownership • Adding free capacity (drives) • Defragmenting free capacity • Placing controllers online or offline • Initializing an array • Reviving an array • Checking redundancy • Deleting an array **Note:** These menu options are only available when a an array is selected.
Logical Drive	The **Logical Drive** menu provides options to perform the following storage management operations on volumes: • Creating logical drives • Changing ownership or segment size • Increasing capacity • Iinitializing, deleting, or renaming a logical drive • Creating a VolumeCopy • Viewing VolumeCopies using the Copy Manager • Creating or disabling a FlashCopy logical drive • Creating a remote mirror • Viewing logical drive properties **Note:** These menu options are only available when a logical drive is selected.
Controller	The **Controlle**r menu displays options to perform the following storage management operations on controllers: • Placing a controller online or offline • Enabling or disabling data transfer (I/O) • Changing the controller mode to active • Changing the preferred loop ID • Modify the IP address, gateway address, or network subnet mask of a controller • Running diagnostics • Viewing controller properties **Note:** These menu options are only available when a drive is selected.

Table 12. The Subsystem Management window menus (continued)

Menu	Definition
Drive	The **Drive** menu contains options to perform the following storage management operations on drives: • Locating a drive • Assigning or unassigning a hot spare • Failing, reconstructing, reviving or initializing a drive • Viewing drive properties **Note:** These menu options are only available when a controller is selected.
Advanced	The **Advanced** menu presents maintenance options which should only be used under the guidance of technical support.
Help	The **Help** menu provides options to perform the following actions: • Display the contents of the Subsystem Management window help • View a reference of all Recovery Guru procedures • View the software version and copyright information

The script editor

Instead of navigating through the GUI interface to perform storage subsystem management functions, a script editor window, as shown in Figure 10 on page 29 is provided for running scripted management commands. Currently, some of the management functions that can be done through the GUI are not implemented through script commands.

```
Tests - F600B4 - Script Editor - CreateLUN.txt                    _ □ ×

File  Edit  Tools  Help

create  logicalDrive  driveCount=4 userLabel="2882_l_R51" RAIDLevel=5
capacity=4GB owner=a;
create  logicalDrive  array=l  userLabel="2882_l_R52" RAIDLevel=5 capacity=4GB
owner=b;

create hostGroup userLabel="MPROGrp";
create host userLabel="MPRO" hostGroup="MPROGrp";
create hostport UserLabel="MPRO1" host="MPRO" identifier="210000e08b04e23d"
hostType=5;
create hostport UserLabel="MPRO2" host="MPRO" identifier="210000e08b012a63"
hostType=5;

───────────────────────────────────────────────────────────────────

Performing syntax check...
Syntax check complete.
```

SJ001138

Figure 10. The script editor window

Important: Use caution when running the commands in the script window because
the script editor does not prompt for confirmation on operations that are destructive
such as the **Delete arrays, Reset Storage Subsystem configuration** commands.

Not all script commands are implemented in all versions of the firmware. The earlier
the firmware version, the smaller the set of script commands. For more information
about script commands and firmware versions, see the online help and the
scriptcommands.pdf file in the DOCS folder of the FAStT Storage Manager CD.

For a list of available commands and their syntax, see the online Command
Reference help.

Using the script editor
Perform the following steps to open the script editor:

1. Select a storage subsystem in the Device Tree view or from the Device table.

2. Select **Tools → Execute Script**.

3. The script editor opens. The Script view and the Output view are presented in
   the window.

   • The Script view provides an area for inputting and editing script commands.
     The Script view supports the following editing key strokes:

     – Cntl+A: To select everything in the window

     – Cntl+C: To copy the marked text in the window into a Windows clipboard
       buffer

     – Cntl+V: To paste the text from the Windows clipboard buffer into the
       window

– Cntl+X: To delete (cut) the marked text in the window
  – Cntl+Home: To go to the top of the script window
  – Cntl+End: To go to bottom of the script window
- The Output view displays the results of the operations.

A splitter bar divides the window between the Script view and the Output view. Drag the splitter bar to resize the views.

The following list includes some general guidelines for using the script editor:

- All statements must end with a semicolon (;).
- Each base command and its associated primary and secondary parameters must be separated by a space.
- The script editor is *not* case sensitive.
- Each new statement must begin on a separate line.
- Comments can be added to your scripts to make it easier for you and future users to understand the purpose of the command statements.

Adding comments to a script

The script editor supports the following comment formats:

- Text contained after two forward slashes (//) until an end-of-line character is reached

  For example:

```
//The following command assigns hot spare drives.
set drives [1,2 1,3] hotspare=true;
```

  The comment //The following command assigns hot spare drives. is included for clarification and is not processed by the script editor.

  Important: You must end a comment that begins with // with an end-of-line character, which you insert by pressing the Enter key. If the script engine does not find an end-of-line character in the script after processing a comment, an error message displays and the script fails.

- Text contained between the /* and */ characters

  For example:

```
/* The following command assigns hot spare drives.*/
set drives [1,2 1,3] hotspare=true;
```

  The comment /*The following command assigns hot spare drives.*/ is included for clarification and is not processed by the script Eeditor.

  Important: The comment must start with /* and end with */. If the script engine does not find both a beginning and ending comment notation, an error message displays and the script fails.

The command line interface (SMcli)

You can use the command line interface (SMcli) to perform the following tasks:

- Run scripts on multiple storage systems
- Create batch files
- Run mass operations on multiple storage systems
- Access the script engine directly without using the Enterprise Management window

Using SMcli

Perform the following steps to use the SMcli:

1. Go to the command line shell of your operating system. At the command prompt, type `SMcli`, followed by either the controller name, host-agent name, worldwide name (WWN) or user-supplied name of the specific storage subsystems. The name that you enter depends on your storage subsystem management method:

    - For directly managed subsystems, enter the hostname or IP address of the controller or controllers
    - For host-agent managed subsystems, enter the hostname or IP address of the host

    Note: Some command line shells might not support commands longer than 256 characters. If your command is longer than 256 characters, use a different shell or enter the command into the Storage Manager script editor.

    If you specify a host name, or the IP address, the command line utility verifies that a storage subsystem exists.

    If you specify the user-supplied storage subsystem name or WWN, the utility ensures that a storage subsystem with that name exists at the specified location and can be contacted.

    Notes:

    - You must use the `-n` parameter if more than one host-agent managed storage subsystem is connected to the host. For example:

        `SMcli hostmachine -n sajason`

    - Use the `-w` parameter if you specify the WWN of the storage subsystem. For example:

        `SMcli -w 600a0b800006602d000000003beb684b`

    - You can specify the storage subsystem by its user supplied name only using the `-n` parameter if the storage subsystem is configured in the Enterprise Management window. For example:

        `SMcli -n Storage Subsystem London`

        The name must be unique to the Enterprise Management window.

2. Type one or more commands, for example:

    `-c "<command>;[<command2>;...]"`

    or

    type the name of a script file, for example:

    `-f <scriptfile>.`

    SMcli first verifies the existence and locations of the specified storage subsystems and, if applicable, the script file. Next, it verifies the script command syntax and then runs the commands.

3. Then you can do one of the following actions:

    - Specify the output file, for example:

        `[-o <outputfile>]`

    - Specify the password, for example:

        `[-p <password>]`

    - Run the script only, for example:

[-e]

Note: These arguments are optional.

Command line interface parameters

The command line interface supports the command line parameters that are described in Table 13.

Table 13. Command line parameters

Command line parameter	Action
`<IP address>` or `<hostname>`	Specify an IP address (xx.xx.xx.xx) or host name (of a host-agent or controller) of a storage subsystem that is managed through the host-agent or direct-management method.
`-a`	Add an SNMP trap destination or e-mail alert destination. For example: Use the following command to add an SNMP trap destination: • `-a trap:COMMUNITY,HOST` where — `COMMUNITY` is the SNMP community name set in the NMS configuration file by a network administrator. The default is public. — `HOST` is the IP address or the host name of a station running an SNMP service. At a minimum, this will be the network management station. Use the following command to delete an e-mail alert destination: `-a email:MAILADDRESS` where `MAILADDRESS` is the fully qualified e-mail address to which the alert message should be sent **Important:** There is no space after the colon (`:`) or the comma (`,`).
`-A`	Specify a storage subsystem to add to the management domain. Specify an IP address (xx.xx.xx.xx) for each controller in the storage subsystem. **Important:** If you specify one IP address, the storage subsystem is partially managed. If no IP address is specified, an automatic discovery of storage subsystems that are attached to the local subnet is performed.

Table 13. Command line parameters (continued)

Command line parameter	Action
-c	Specify the list of commands to be performed on the specified storage subsystem. Consider the following usage requirements: • You cannot place multiple -c parameters on the same command line. However, you can include multiple commands after the -c parameter. • Each command must end with a semicolon (;). • In Microsoft Windows environments, the entire command string must be enclosed in double quotation marks ("). Each command must end with a semicolon (;). • In UNIX environments, the entire command string must be enclosed in single quotation marks ('). Each command must end with a semicolon (;). **Note:** Any errors that are encountered when running the list of commands will by default cause the command to stop. Use the **on error continue;** command first in the list of commands to override this situation.
-d	Display the contents of the configuration file in the following format: `<storagearrayname> <hostname> <hostname>` The configuration file lists all known storage subsystems that are currently configured in the Enterprise Management window.
-e	Run the commands only, without performing a syntax check first.
-f	Specify the name of a file containing script engine commands to be performed on the specified storage subsystem. Use the -f parameter in place of the -c parameter. **Note:** Any errors that are encountered when running the list of commands will by default cause the command to stop. Use the **on error continue;** command in the script file to override this situation.
-F	Specify the e-mail address that will send the alerts.
-i	When used with the -d parameter, display the contents of the configuration file in the following format: `<storagearrayname> <IP address> <IP address>`
-m	Specify the IP address or host name of the mail or SNMP server that will send the alerts.
-n	Specify the storage subsystem name on which you want to perform the script commands. This name is optional when a `<hostname or IP address>` is used. However, if you are managing the storage subsystem using the host-agent management method, you must use the -n parameter if more than one storage subsystem is connected to the host at the specified address. This name is required when the `<hostname or IP address>` is not used. However, the storage subsystem name must be configured for use in the Enterprise Management window and must not be a duplicate of any other configured storage subsystem name.

Table 13. Command line parameters (continued)

Command line parameter	Action
-o	Specify a file name for all output text from the script engine. If this parameter is not used, the output will go to a standard output device.
-p	Specify the password for the storage subsystem on which you want to perform a command script. A password is not necessary under the following curcumstances: • A password has not been set on the storage subsystem. • The password is specified with the **use password** command in the script file with the -f parameter. • You specify the password with the **use password** command using the -c parameter.
-s	Display the alert settings for the storage subsystems that are currently configured in the Enterprise Management window.
-w	Specify the WWN of the device on which you want to perform script commands. When used in conjunction with -d, displays the WWN of the storage subsystems contained in the configuration files.
-x	Delete an SNMP trap destination or e-mail alert destination. Use the following command to delete an SNMP trap destination: -x trap:COMMUNITY, HOST where • COMMUNITY is the SNMP community name • HOST is the IP address or the host name of a station running an SNMP service. Use the following command to delete an e-mail alert destination: -x email:MAILADDRESS where MAILADDRESS is the fully qualified e-mail address to which the alert message should no longer be sent.
-?	Display usage information

Notes:

• All statements must end with a semi-colon (;).
• Separate each base command and any parameters with a space.
• Separate each parameter and its parameter value with an equal sign (=).
• The SMcli is not case-sensitive. You can enter any combination of upper and lowercase letters. The usage shown in the examples in the section "SMcli examples" on page 36 follows the convention of having a capital letter start the second word of a parameter.
• For a list of supported commands and their syntax, see the Enterprise Management window online help. The online help contains commands that are current with the latest version of the storage management software.

  Some of the commands might not be supported if you are managing storage subsystems running firmware for previous releases. See the Firmware

Compatibility List in the Enterprise Management window online help for a complete list of commands and the firmware levels on which they are supported.

Usage and formatting requirements

SMcli has the following usage and formatting requirements:

For all operating systems:

- If you invoke SMcli with no arguments or with an unrecognized parameter, usage information is displayed.
- Arguments following the `-n`, `-o`, `-f`, and `-p` parameters that contain a space or a special character (<, >, ', !, *, for example) must be enclosed in single quotation marks (') or double quotation marks ("), depending on your operating system.

  Use single quotation marks (') if you are using a UNIX operating system and use double quotation marks (") if you are using a Windows operating system. For examples of the differences in expressing arguments for UNIX and Windows operating systems, see "SMcli examples" on page 36.
- Arguments following the `-n`, `-o`, `-f`, and `-p` parameters that contain a single quotation mark (') must be enclosed in double quotation marks (").

  Note: If you invoke SMcli and specify a storage subsystem, but do not specify the commands or script file to run, SMcli runs in interactive mode. This allows you to specify the commands interactively. Use Ctrl+D to stop SMcli.

For Microsoft Windows operating systems only:

- Insert a backslash (\) before each double quotation marks (") when the double quotation marks are used as part of a name or command syntax. For example:

  `-c "set storageSubsystem userLabel=\"string\";"`
- Insert three backslashes (\\\) in front of the double quotation marks (") to display the backslash when used with the `-n`, `-o`, `-f`, or `-p` parameter. For example, to specify storage subsystem named `Jason\`, type:

  `-n "Jason\\\"`
- Insert five backslashes (\\\\\) in front of the double quotation marks (") to use the backslash character as part of the literal command string. For example, to change the name of the storage subsystem to `Jason\`, type the following command:

  `-c "set storageSubsystem userLabel="Jason\\\\\";"`
- Insert a caret (^) before each special script character (^, &, |, <, >) when that character is used with the `-n`, `-o`, `-f`, and `-p` parameters. For example, to specify storage subsystem "CLI&CLIENT", type:

  `-n "CLI^&CLIENT"`
- Insert three carets (^^^) before each special script character when it is used within a literal script command string. For example, to change the name of the storage subsystem to `Finance&Payroll`, type the following command:

  `-c "set storageSubsystem userLabel=\"Finance^^^&payroll\";"`

See the appropriate operating system documentation for a list special script characters.

SMcli examples

Following are examples of how you can use the SMcli to access and run script engine commands.

Note: The usage of the -c and the -p parameters varies depending on your operating system.

- For Microsoft Windows systems, the -c and the -p parameters must be enclosed in double quotation marks (").
- For UNIX systems, the -c and the -p parameter strings must be enclosed in single quotation marks (').

1. Rename "Payroll Array" to "Finance Array" using the hostname ICTSANT.

   For Windows systems:

   ```
    SMcli ICTSANT -n "Payroll Array" -c "set storageSubsystem
   userlabel=\"Finance Array\";"
   ```

   For UNIX systems:

   ```
   SMcli ICTSANT -n 'Payroll Array' -c 'set storageSubsystem
   userlabel="Finance Array";'
   ```

2. In the storage subsystem with controller names "finance 1" and "finance 2," use the password `Test Array` to do the following:

   - Delete the logical drive named "Stocks & Bonds".
   - Create a new logical drive named "Finance".
   - Show the health status of the storage subsystem, which is managed using the direct management method.

   For Windows systems:

   ```
   SMcli finance1 finance2 -c "use password"TestArray";
   delete logicalDrive[\"Stocks^^^&Bonds\"];
   create logicalDrive driveCount[3] RAIDLevel=3 capacity=10GB userLabel=\"Finance\";
   show storageSubsystem healthStatus;"
   ```

   For UNIX systems:

   ```
   SMcli finance1 finance2 -c 'use password "TestArray";
   delete logicalDrive ["Stocks&Bonds"];
   create logicalDrive driveCount[3] RAIDLevel=3 capacity=10GB userLabel="Finance";
   show storageSubsystem healthStatus;'
   ```

3. Run the commands that are in the script file named scriptfile.scr in the storage subsystem named "Example" without performing a syntax check.

   For both Windows and UNIX systems:

   ```
   SMcli -n Example -f scriptfile.scr -e
   ```

4. Run the commands found in the script file named scriptfile.scr on the storage subsystem named "Example" . Use "My Array" as the password and direct all output to output.txt.

   For Windows systems:

   ```
   SMcli -n Example -f scriptfile.scr -p "My Array" -o output.txt
   ```

For UNIX systems:

```
SMcli -n Example -f scriptfile.scr -p 'My Array' -o output.txt
```

5. Display all storage subsystems that are currently configured in the Enterprise Management window (configuration file), using <IP address> format instead of <hostname> format.

   For Windows and UNIX systems:

```
SMcli -d —i
```

Chapter 2. Protecting your data

When you configure 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 24 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 controller firmware configurations are:
- Support for different logical-drive RAID levels
- Orthogonal RAID striping 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

Logical drives

The storage management software identifies several distinct logical drives. The following list describes each type of logical drive.

Standard logical drive
> A standard logical drive is a logical structure that is created on a storage subsystem for data storage. Use the Create Logical Drive wizard to create a standard logical drive. Only standard logical drives are created if neither FlashCopy nor the Remote Mirror Option features are enabled. Standard logical drives are also used with creating FlashCopy logical drives and Remote Mirror logical drives.

FlashCopy logical drive
> A FlashCopy logical drive is a point-in-time image of a standard logical drive. A FlashCopy logical drive is the logical equivalent of a complete physical copy, but you create it much more quickly and it requires less disk space. The logical drive from which you are creating the FlashCopy logical drive, called the base logical drive, must be a standard logical drive in your storage subsystem. For more information about FlashCopy logical drives, see "FlashCopy" on page 50.

FlashCopy repository logical drive
> A FlashCopy repository logical drive is a special logical drive in the storage

subsystem that is created as a resource for a FlashCopy logical drive. A FlashCopy repository logical drive contains FlashCopy logical drive metadata and copy-on-write data for a particular FlashCopy logical drive. For more information about FlashCopy repository logical drives, see "FlashCopy" on page 50.

Primary logical drive
A primary logical drive is a standard logical drive in a mirror relationship that accepts host I/O and stores application data. When the mirror relationship is first created, data from the primary logical drive is copied in its entirety to the associated secondary logical drive. For more information about primary logical drives, see "Remote Mirror option" on page 52.

Secondary logical drive
A secondary logical drive is a standard logical drive in a mirror relationship that maintains a mirror (or copy) of the data from its associated primary logical drive. The secondary logical drive remains unavailable to host applications while mirroring is active. In the event of a disaster or catastrophic failure of the primary site, you can promote the secondary logical drive to a primary role. For more information about secondary logical drives, see "Remote Mirror option" on page 52.

Mirror repository logical drive
A mirror repository logical drive is a special logical drive that is created as a resource for each controller in both the local and remote storage subsystem. The controller stores mirroring information on the mirror repository logical drive which includes information about remote writes that are not yet complete. The controller can use the mirrored information to recover from controller resets and accidental powering-down of storage subsystems. For more information about mirror repository logical drives, see "Remote Mirror option" on page 52.

Source logical drive
A source logical drive is a standard logical drive that contains the data that, through a VolumeCopy operation, will be copied to another logical drive, which is known as the target logical drive. A source logical drive can be either a standard logical drive, a FlashCopy logical drive, the base logical drive of a FlashCopy logical drive, or a primary logical drive of a mirrored pair. For more information about source logical drives, see "VolumeCopy" on page 51.

Target logical drive
A target logical drive is a standard logical drive to which the data on the source logical drive is copied during a VolumeCopy operation. When a logical drive is selected as a target logical drive, any existing data on the logical drive is completely overwritten and the logical drive automatically becomes read-only after the copy operation has completed, to protect it from host write access. After the logical drive copy completes, you can use the Copy Manager to disable the Read-Only attribute for the target logical drive. For more information about target logical drives, see "VolumeCopy" on page 51.

Dynamic Logical Drive Expansion

Attention: Increasing the capacity of a standard logical drive is only supported on certain operating systems. If you increase the logical drive capacity on a host operating system that is unsupported, the expanded capacity will be unusable, and you cannot restore the original logical drive capacity. For information about supported operating systems, see Increase Logical Drive Capacity: Additional Instructions in the Storage Subsystem Management window online help.

Dynamic Logical Drive Expansion (DVE) is a modification operation that you use to increase the capacity of standard or FlashCopy repository logical drives. You can increase the capacity by using any free capacity available on the array of the standard or FlashCopy repository logical drive.

Data is accessible on arrays, logical drives, and disk drives throughout the entire modification operation.

During the modification operation, the logical drive for which the capacity is being increased shows the following three factors:
- A status of `Operation in Progress`
- The original logical drive capacity
- The total capacity being added

After the capacity increase completes, the expanded capacity of the logical drive displays, and the final capacity for the Free Capacity node that is involved shows a reduction in capacity. If you use all of the free capacity to increase the logical drive size, then the Free Capacity node that is involved is removed from the Logical View.

You cannot increase the storage capacity of a logical drive if any of the following conditions apply:
- One or more hot spare drives are in use in the logical drive
- The logical drive has a non-optimal status
- Any logical drive in the array is in any state of modification
- The controller that owns this logical drive is in the process of adding capacity to another logical drive. (Each controller can add capacity to only one logical drive at a time.)
- No free capacity exists in the array
- No unconfigured capacity (in the form of drives) is available to add to the array

Attention: An increase in storage capacity for FlashCopy repository logical drives completes if a warning is received that the FlashCopy repository logical drive is in danger of becoming full. Increasing the capacity of a FlashCopy repository logical drive does not increase the capacity of the associated FlashCopy logical drive. The capacity of the FlashCopy logical drive is always based on the capacity of the base logical drive when the FlashCopy logical drive is created.

For more information, see Learn About Increasing the Capacity of a Logical Drive on the **Learn More** tab in the Storage Subsystem Management online help window.

Arrays

An array is a set of drives that the controller logically groups together to provide one or more logical drives to an application host. When you create a logical drive

from unconfigured capacity, the array and the logical drive are created at the same time. When you create a logical drive from free capacity, an additional logical drive is created on an existing array.

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

Dynamic Capacity Expansion

Dynamic Capacity Expansion (DCE) is a modification operation that you use to increase the available free capacity on an array. The increase in capacity is achieved by selecting unassigned drives to be added to the array. After the capacity expansion is completed, additional free capacity is available on the array for creation of other logical drives. The additional free capacity can then be used to perform a DVE on a standard or FlashCopy repository logical drive.

This modification operation is considered to be "dynamic" because you have the ability to continually access data on arrays, logical drives, and disk drives throughout the entire operation. For more information, see Learn About Increasing the Capacity of a Logical Drive on the **Learn More** tab of the online help.

Fibre-channel I/O data path failover support

I/O data path protection to redundant controllers in a FAStT storage subsystem is accomplished with either the Auto-Logical Drive Transfer (ADT) feature, a host multipath driver, or both. A multipath driver, such as the redundant disk array controller (RDAC), is an I/O path failover driver that is installed on host computers that access the storage subsystem. This section describes ADT and other operating-system specific failover protection features.

Auto-Logical Drive Transfer feature

The Auto-Logical Drive Transfer (ADT) feature is a built-in feature of the controller firmware that allows logical drive-level failover protection rather than controller-level failover protection. For controller firmware versions 05.2x.xx.xx and higher, the ADT feature is automatically disabled or enabled depending on the type of host ports in the host partition to which you mapped the logical drives. It is disabled by default for Microsoft Windows, IBM AIX, and Sun Solaris operating systems. It is enabled by default for Linux, Novell Netware, and HP-UX operating systems.

Redundant array of independent disks

Redundant array of independent disks (RAID) is available on all operating systems and relies on a series of configurations, called *levels*, to determine how user and redundancy data is written and retrieved from the drives. The FAStT Storage Manager software offers four RAID level configurations:

- RAID-0
- RAID-1
- RAID-3
- RAID-5

Each level provides different performance and protection features.

RAID-1, RAID-3, and RAID-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 that is 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. RAID-1 uses mirroring for redundancy. RAID-3 and RAID-5 use redundancy information, sometimes called *parity*, that is constructed from the data bytes and striped along with the data on each disk.

Table 14 describes the RAID level configurations that are available with the FAStT Storage Manager Version 8.4 software.

Table 14. RAID level configurations

RAID level	Description
RAID-0	RAID-0 offers simplicity, but does not provide data redundancy. A RAID-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	RAID-1 uses disk mirroring to make an exact copy of data from one drive to another drive. If one drive fails in a RAID-1 array, the mirrored drive takes over. RAID-1 is costly in terms of capacity. One-half of the drives are used for redundant data.
RAID-3	RAID-3 requires one dedicated disk in the logical drive to hold redundancy information. User data is striped across the remaining drives. RAID-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.
RAID-5	RAID-5 spreads redundancy information across all of the drives in the logical drive. RAID-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.

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 that is reserved for redundancy data. The amount of capacity that is needed for redundancy depends on the RAID level that is used.

A redundancy check can be performed using the Check Redundancy menu option. The redundancy check performs one of the following actions:
- Scans the blocks in a RAID-3 or RAID-5 logical drive and checks the redundancy information for each block
- Compares data blocks on RAID-1 mirrored drives

Redundant disk array controller

The redundant disk array controller (RDAC) driver manages the fibre-channel I/O path failover process for storage subsystems in Microsoft Windows NT 4, Windows 2000, Windows Server 2003, IBM AIX, and Sun Solaris environments 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. 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.

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

In Storage Manager versions prior to version 8.xx, there is a single path limitation for the Microsoft Windows and Sun Solaris RDAC. This means that in a given server, the RDAC driver must see only one path from the port of the HBA to the storage controller in the storage subsystem. In Storage Manager version 8.xx and higher, this limitation has been relaxed to four paths.

Note: If there is more than one path from the host to the storage server, RDAC sends the I/O requests down each of the paths using a round robin schedule.

Operating system specific failover protection

The following list indicates which features are available for failover protection for the specified operating systems.

Microsoft Windows
> The RDAC driver is provided in the FAStT storage management software package.

Novell NetWare
> For Storage Manager 8.4 and later, use the Novell native multi-path driver and the LSIMPE.CDM file for failover protection. For Storage Manager 8.3 and earlier, the IBM Storage Area Network (IBMSAN) driver is provided with the FAStT storage management software package. The IBMSAN 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 FAStT HBA driver. In case of a failure along the I/O path, the driver sends the I/O requests through the alternate path and ADT moves the logical drive within the FAStT Storage Server.
>
> For more information on failover protection, see the *IBM TotalStorage FAStT Storage Manager 8.4 Installation and User's Guide for Intel-based Operating System Environments*.
>
> For the current readme file, IBMSAN driver, and installation instructions, go to the following Web site:
>
> www.ibm.com/pc/support

Linux A failover version of the fibre-channel HBA driver is available for failover protection.

In addition, you must use a fibre-channel diagnostic program called FAStT-MSJ to manually assign logical drives to a preferred path between the FAStT storage controllers and the HBAs in the Linux host. This program installs a QLremote agent which must be running during the path failover configuration task.

AIX and Sun Solaris
The RDAC driver is provided in the FAStT storage management software package.

HP-UX
The standard HP-UX operating system installation includes a fibre-channel HBA driver that provides built-in multipath failover support.

Default settings for failover protection

The storage management software uses the following default settings, based on the host type:

- Multipath driver software on the host or hosts and ADT enabled on the storage subsystem
- Multipath driver software on the host or hosts and ADT disabled on the storage subsystem
- No multipath driver software on the host or hosts and ADT enabled on the storage subsystem (no failover)

Note: If you want to change the default ADT settings, contact technical support.

Multipath driver software with ADT enabled on the storage subsystem
This is the normal configuration setting for Novell Netware, Linux, and Hewlett Packard HP-UX systems.

Two active controllers are located in a storage subsystem. When you create a logical drive, you assign one of the two active controllers to own the logical drive (called *preferred controller ownership*) and to control the I/O between the logical drive and the application host along the I/O path. The preferred controller normally receives the I/O requests from the logical drive. If a problem along the data path (such as a component failure) causes an I/O to fail, the multipath driver issues the I/O to the alternate controller.

When ADT is enabled and used with a host multipath driver, it helps ensure an I/O data path is available for the storage subsystem logical drives. The ADT feature changes the ownership of the logical drive that is receiving the I/O to the alternate controller. After the I/O data path problem is corrected, the preferred controller automatically reestablishes ownership of the logical drive as soon as the multipath driver detects that the path is normal again.

Multipath driver software with ADT disabled on the storage subsystem
This is the configuration setting for Microsoft Windows, IBM AIX, and Sun Solaris systems.

When ADT is disabled, the I/O data path is still protected as long as you use a multipath driver. However, when an I/O request is sent to an individual logical drive and a problem occurs along the data path to its preferred controller, all logical drives on the preferred controller are transferred to the alternate controller. In addition, after the I/O data path problem is corrected, the preferred controller does not automatically reestablish ownership of the logical drive. You must open a storage management window and perform the Redistribute Logical Drives task.

No multipath driver software with ADT enabled on the storage subsystem (no failover protection)

> **Note:** This setting is not supported.
> The FAStT storage subsystems in this scenario have no failover protection. A pair of active controllers might still be located in a storage subsystem and each logical drive on the storage subsystem might be assigned a preferred owner. However, logical drives do not move to the alternate controller because there is no multipath driver installed. When a component in the I/O path, such as a cable or the controller itself, fails, I/O operations cannot get through to the storage subsystem. The component failure must be corrected before I/O operations can resume. You must switch logical drives to the alternate controller in the pair manually.
>
> **Note:** Hosts that use operating systems without failover capability should be connected to the storage subsystem so that each host adapter has only one path to the controller.

Protecting data in the controller cache memory

Write caching enables the controller cache memory to store write operations from the host computer, which improves 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 that is 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 controller completes all outstanding write operations.

Note: You can enable the write-cache mirroring parameter for each logical drive but when write-cache mirroring is enabled, half of the total cache size in each controller is reserved for mirroring the cache data from the other controller.

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, the default start and stop settings for a logical drive are 80% and 20% of the total cache size, respectively. With these settings, the controller starts flushing the cache data when the cache reaches 80% full and stops flushing cache data when the cache is flushed down to 20% full. For maximum data safety, you can choose low start and stop percentages, for example, a start setting of 25% and a stop setting of 0%. However, these low start and stop settings increase the chance that data that is 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. There are batteries in the controller enclosure that protect the cache against power outages. The controller battery backup CRU change interval is three years from the date that the backup battery CRU was installed. 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 of the battery (in days) so that you know when it is time to replace the battery.

Note: For the FAStT200, FAStT600 and FAStT600 Turbo storage servers, the battery CRU is located inside each controller CRU.

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 the batteries in the controller enclosure are removed.

Attention: For maximum data integrity, 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 that contains no data and that acts as a standby in case a drive fails in a RAID-1, RAID-3, or RAID-5 array. If the logical drive in a RAID-1, RAID-3, or RAID-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*.

For FAStT controllers with 04.01.xx.xx and 05.xx.xx.xx firmware, the hot-spare drive selection algorithm searches for a suitable hot-spare drive within the same external drive enclosure as the failed drive before looking for the hot-spare drive in other enclosures. For maximum data protection, you should use only the largest capacity drives for hot-spare drives in mixed capacity hard drive configurations.

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 media scan 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.

The background media scan can be enabled with or without checking for the logical drive redundancy data. For RAID-3 or RAID-5 logical drives, a redundancy data check scans the data blocks, calculates the redundancy data, and compares it to the read redundancy information for each block. It then repairs any redundancy errors, if required. For a RAID-1 logical drive, a redundancy data check compares data blocks on mirrored drives and corrects any data inconsistencies. When

redundancy checking is not enabled, the FAStT storage server still scans all blocks in the logical drives, including the redundancy blocks but it does not check for the accuracy of the redundancy data.

Note: The media scan must be enabled for the entire storage subsystem and enabled on each logical drive within the storage subsystem.

When enabled, the media scan runs on all logical drives in the storage subsystem that meet the following conditions:

- The logical drive is in an optimal status
- There are no modification operations in progress
- The Media Scan parameter is enabled on the logical drive **Change Media Scan Settings** dialog box

Because additional I/O reads are generated for media scanning, there might be a slight performance impact depending on the following factors:

- The amount of configured storage capacity in the FAStT storage subsystem
- The configured scan duration for the media scan operations
- The status of the redundancy check option (enabled or disabled)

For controllers with 4.01.xx.xx and 05.xx.xx.xx firmware, the background media scan operation does not scan hot-spare or unused optimal hard drives (those that are not part of a defined logical drive) in a FAStT storage server configuration. To perform a media scan on hot-spare or unused optimal hard drives, you must convert them to logical drives at certain scheduled intervals and then revert them back to their hot-spare or unused states after you scan them.

Errors reported by a media scan

The media scan process runs continuously in the background when it is enabled. Every time a scan cycle (that is, a media scan of all logical drives in a storage subsystem) completes, it restarts immediately. The media scan process discovers any errors and reports them to the storage subsystem event log (MEL). Table 15 lists the errors that are discovered during a media scan.

Table 15. Errors discovered during a media scan

Error	Description and result
Unrecovered media error	The drive could not read the data on its first attempt, or on any subsequent attempts. For logical drives or arrays with redundancy protection (RAID-1, RAID-3 and RAID-5), data is reconstructed, rewritten to the drive, and verified. The error is reported to the event log. For logical drives or arrays without redundancy protection (RAID-0 and degraded RAID-1, RAID-3, and RAID-5 logical drives), the error is not corrected but is reported to the event log.
Recovered media error	The drive could not read the requested data on its first attempt but succeeded on a subsequent attempt. The data is rewritten to the drive and verified. The error is reported to the event log. **Note:** Media scan makes three attempts to read the bad blocks.

Table 15. Errors discovered during a media scan (continued)

Error	Description and result
Redundancy mismatches	Redundancy errors are found. The first 10 redundancy mismatches that are found on a logical drive are reported to the event log. **Note:** This error could occur only when the optional redundancy checkbox is enabled, when the media scan feature is enabled, and the logical drive or array is not RAID-0.
Unfixable error	The data could not be read and parity or redundancy information could not be used to regenerate it. For example, redundancy information cannot be used to reconstruct data on a degraded logical drive. The error is reported to the event log.

Media scan settings

To maximize the protection and minimize the I/O performance impact, the FAStT storage subsystem is shipped with the following default media scan settings:

- The media scan option is enabled for all logical drives in the storage subsystem. Therefore, every time a logical drive is created, it is created with the media scan option enabled. If you want to disable media scanning, you must disable it manually for each logical drive.
- The media scan duration is set to 30 days. This is the time in which the FAStT controllers must complete the media scan of a logical drive. The controller uses the media scan duration, with the information about which logical drives must be scanned, to determine a constant rate at which to perform the media scan activities. The media scan duration is maintained regardless of host I/O activity.

  Thirty days is the maximum duration setting. You must manually change this value if you want to scan the media more frequently. This setting is applied to all logical drives in the storage subsystem. For example, you cannot set the media scan duration for one logical drive at two days and the others logical drives at 30 days.
- The redundancy check option is not enabled. You must manually set this option for each of the logical drives that you want to have redundancy data checked.

Note: Changes to the media settings will not be in effect until the current media scan cycle completes or both controllers are reset.

To change the media scan settings for the entire storage subsystem, perform the following steps:

1. Select the storage subsystem entry in the Logical/Physical view of the Subsystem Management window.
2. Click **Storage Subsystem → Change → Media Scan Settings**.

To change the media scan settings for a given logical drive, perform the following steps:

1. Select the logical drive entry in the Logical/Physical view of the Subsystem Management window.
2. Click **Logical Drive → Change → Media Scan Settings**.

Copy services and the FAStT Storage Server

FAStT Storage Manager 8.4 supports the following copy service features, which are available for purchase separately from IBM or an IBM Business Partner:

FlashCopy

> The FlashCopy premium feature supports creating and managing of FlashCopy logical drives. A FlashCopy is the logical equivalent of a complete physical copy, but is created more quickly and requires less disk space. It is host addressable, so you can perform backups using FlashCopy while the base logical drive is online and user-accessible. When the backup completes, you can delete the FlashCopy logical drive or save it for reuse.

VolumeCopy

> The VolumeCopy premium feature is a new feature that is supported on firmware version 5.4x.xx.xx only. It is a firmware-based mechanism for replicating data within a storage array and is used with FlashCopy. This feature is designed as a system management tool for tasks such as relocating data to other drives for hardware upgrades or performance management, data backup, and data restoration.

Remote Mirror option

> The Remote Mirror option premium feature provides online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Mirror option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations.

Note: Table 16 lists the restrictions that apply to the copy service features that are supported on the FAStT200, FAStT600, and the FAStT600 Turbo storage servers.

Table 16. Restrictions to premium feature support on the FAStT200, FAStT600 and FAStT600 Turbo storage servers

Storage server	Features not supported on firmware version 5.3x.xx.xx (See note)	Features not supported on firmware version 5.4x.xx.xx
FAStT200	Remote Mirror option	Remote Mirror option VolumeCopy
FAStT600	Remote Mirror option FlashCopy	Remote Mirror option
FAStT600 Turbo	Remote Mirror option FlashCopy	Remote Mirror option

Note: The VolumeCopy feature is not available on Storage Manager 8.3 and earlier.

FlashCopy

Use FlashCopy to create and manage FlashCopy logical drives. A FlashCopy logical drive is a point-in-time image of a standard logical drive in your storage subsystem. The logical drive that is copied is called a *base logical drive*.

When you make a FlashCopy, the controller suspends writes to the base logical drive for a few seconds while it creates a FlashCopy *repository logical drive*. This is a physical logical drive where FlashCopy metadata and copy-on-write data are stored.

Flashcopy is implemented using a "copy-on-write" scheme:
- FlashCopy logical drive read data comes from the base logical drive if the read data blocks were not modified. Otherwise, it comes from the repository logical drive.
- Base logical drive read data comes from the base logical drive.
- Writes to the base logical drive cause the data in the affected blocks be copied to the repository logical drive if this is the first write to the data blocks. After the original data is copied to the repository logical drive, subsequent writes to the same data blocks will not cause any additional data to be copied to the repository logical drive.

   Note: The data is copied to the repository logical drive sequentially.

You can create up to four FlashCopy logical drives of a base logical drive and then write data to the FlashCopy logical drives to perform testing and analysis. For example, before upgrading a database management system, you can use FlashCopy logical drives to test different configurations. You can disable the FlashCopy when you are finished with it, for example after a backup completes. Then you can re-create the FlashCopy the next time you do a backup and reuse the same FlashCopy repository logical drive.

For operating-system specific information and instructions for using FlashCopy, see the *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide* or the FlashCopy online help.

VolumeCopy

The VolumeCopy feature is a premium feature that comes with the FAStT Storage Manager 8.4 software and is enabled by purchasing a premium feature key. VolumeCopy is used with FlashCopy and, therefore, it can be purchased together with FlashCopy as a single copy service option, or at a later time as an enhancement to FlashCopy. The VolumeCopy feature is a firmware-based mechanism that is used to copy data from one logical drive (the source logical drive) to another logical drive (the target logical drive) in a single storage subsystem. This feature can be used to perform the following tasks:
- Copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives
- Back up data
- Restore FlashCopy logical drive data to the base logical drive

This feature includes a Create Copy wizard that you can use to create a logical drive copy, and a Copy Manager that you can use to monitor logical drive copies after they have been created.

Copying data for greater access

As your storage requirements for a logical drive change, you can use the VolumeCopy feature to copy data to a logical drive in an array that uses larger capacity disk drives within the same storage subsystem. This provides an opportunity to move data to larger drives (for example, 73 GB to 146 GB), change to drives with a higher data transfer rate (for example, 1 Gbps to 2 Gbps), or change to drives that use new technologies for higher performance.

Backing up data

The VolumeCopy feature allows you to create a backup of a logical drive by copying data from one logical drive to another logical drive in the same storage subsystem. The target logical drive can be used as a backup for the source logical drive, for system testing, or to back up to another device, such as a tape drive.

Restoring FlashCopy logical drive data to the base logical drive

If you need to restore data to the base logical drive from its associated FlashCopy logical drive, the VolumeCopy feature can be used to copy the data from the FlashCopy logical drive to the base logical drive. You can create a logical drive copy of the data on the FlashCopy logical drive, then copy the data to the base logical drive.

Attention: If the logical drive that you want to copy is used in a production environment, the FlashCopy feature must be enabled. A FlashCopy of the logical drive must be created and then specified as the VolumeCopy source logical drive, instead of using the actual logical drive itself. This requirement allows the original logical drive to continue to be accessible during the VolumeCopy operation.

For more information about VolumeCopy, see the *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide*.

Remote Mirror option

The following information is an overview of the Remote Mirror option. For more detailed information about the topic, see the *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide*.

The maximum number of storage subsystems that can participate in a remote mirror configuration is two. The two storage subsystems are called *primary* and *secondary* storage subsystems or *local* and *remote* storage subsystems. These names are used interchangeably to describe remote mirror setups or concepts. The names do not refer to the location of storage subsystems or to the role that storage subsystems have in a remote mirror relationship.

Primary and secondary logical drives on a remote mirror setup

When you create a remote mirror, a mirrored logical drive pair is defined and consists of a primary logical drive at the primary storage subsystem, and a secondary logical drive at a secondary storage subsystem. A standard logical drive might only be defined in one mirrored logical drive pair. The maximum number of supported mirrored logical drive pairs is determined by the storage subsystem model.

The primary and secondary role in a remote mirror setup is implemented at the logical drive level instead of at the storage subsystem level. All logical drives that participate in a remote mirror relationship on a storage subsystem can be in either a primary or secondary role only. The storage subsystem can also have a combination of logical drives in a primary role and logical drives in a secondary role. Whether the logical drive is in a primary or secondary role, it counts towards the maximum number of mirror logical drive pairs that can be defined in a storage subsystem.

With controller firmware 05.30.xx.xx. and earlier, there is a limit to how many logical drives that you can create in a single storage subsystem. When the Remote Mirror Option is enabled, the total number of logical drives that are supported for each storage subsystem is reduced from the number of logical drives that you would have without the Remote Mirror Option enabled.

Primary logical drives: The primary logical drive is the drive that accepts host computer I/O operations and stores program data. When the mirror relationship is first created, data from the primary logical drive is copied (becomes a mirror image) in its entirety to the secondary logical drive. This process is known as a *full synchronization* and is directed by the controller owner of the primary logical drive. During a full synchronization, the primary logical drive remains fully accessible for all normal I/O operations.

When a write request is made to the primary logical drive, the controller owner of the primary logical drive also initiates a remote write request to the secondary drive. The write request does not complete until both write requests are performed. This additional write request keeps the data on the two logical drives in the remote mirror relationship synchronized. Whenever the data on the primary drive and the secondary drive becomes unsynchronized, the controller owner of the primary drive initiates a full synchronization.

Secondary logical drives: The secondary logical drive stores the data that is copied from the primary logical drive associated with it. The controller owner of the secondary logical drive receives remote writes from the controller owner of the primary logical drive and does not accept host computer read or write requests.

The secondary logical drive is unavailable to host computer programs while the mirroring operation is performed. In the event of a disaster or unrecoverable error of the primary storage subsystem, a role reversal is performed to promote the secondary logical drive to the primary logical drive. Host computers are then able to access the newly-promoted logical drive and normal operations can continue.

Mirror repository logical drives: A mirror repository logical drive is a special logical drive in the storage subsystem. It is created as a resource for the controller owner of the primary logical drive in a remote mirror. The controller stores mirrored information on this logical drive, including information about remote writes that are not yet written to the secondary logical drive. The controller can use this information to recover from controller resets and accidental powering-down of storage subsystems.

When you activate the Remote Mirror Option on the storage subsystem, the system creates two mirror repository logical drives, one for each controller in the storage subsystem. An individual mirror repository logical drive is not needed for each mirror logical drive pair.

When you create the mirror repository logical drives, you specify their location. You can either use existing free capacity or you can create an array for the logical drives from unconfigured capacity and then specify the RAID level.

Because of the critical nature of the data that is stored, the RAID level of mirror repository logical drives must be non-zero. The required size of each logical drive is 4 MB for each mirror repository logical drive.

Mirror relationships: Before you define a mirror relationship, the Remote Mirror Option must be enabled on both the primary and secondary storage subsystems. A secondary standard logical drive candidate must be created on the secondary storage subsystem if one does not already exist. It must be a standard logical drive and at least the same size as or larger than the primary logical drive.

When secondary logical drive candidates are available, you can define a mirror relationship in the storage management software by identifying the storage subsystem that contains the primary logical drive and the storage subsystem that contains the secondary logical drive.

When you set up the mirror relationship, a full synchronization occurs as data from the primary logical drive is copied in its entirety to the secondary logical drive.

Data replication

Data replication between the primary logical drive and the secondary logical drive is managed by the controllers and is transparent to host computers and programs.

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

After the host computer write request is written to the primary logical drive and the data is successfully copied to the secondary logical drive, the controller removes the log record on the mirror repository logical drive and sends an I/O completion indication back to the host computer.

Because the controller does not send the I/O completion to the host computer until the data is copied to both the primary and secondary logical drives, this mirroring operation is called *synchronous*.

When a read request is received from a host computer, the controller owner of the primary logical drive manages the request. No communication takes place between the primary and secondary storage subsystems.

Supported configurations

There are three configurations that support the Remote Mirror option when you have FAStT Storage Manager Version 8.4 software installed. The level of redundancy is determined by the type of configuration that you choose. The following configurations support the Remote Mirror option:

Highest Availability Campus configuration
>   The Highest Availability Campus configuration is fully redundant, and includes two storage subsystems and four fibre-channel switches that are connected with a fibre-channel fabric. The primary storage subsystem and secondary storage subsystem have a maximum connection distance of up to 10 Km (6.25 mi).

>   **Note:** For performance and high availability configuration, use two 10 Km (6.25 mi) interswitch links (ISL).

Campus configuration
>   The Campus configuration, a lower-cost configuration, consists of two storage subsystems and two fibre-channel switches that are connected with a fibre-channel fabric. The primary storage subsystem and secondary storage subsystem have a maximum connection distance of up to 10 Km (6.25 mi).

Intra-site configuration
>   The Intra-site configuration, the lowest-cost configuration, consists of two storage subsystems and two fibre-channel switches that are connected with a fibre-channel fabric. The primary storage subsystem and secondary

storage subsystem might have a maximum connection distance of up to 500 m (0.32 mi), and can be located in the same building or in the same room.

If these configurations do not meet your remote mirror configuration requirement, contact your IBM technical support representative. For more information and detailed instructions about setting up the configurations, see the *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide.*

Managing Persistent Reservations

Attention: The Persistent Reservations option should be used only with guidance from an IBM technical-support representative.

The Persistent Reservations option enables you to view and clear volume reservations and associated registrations. Persistent reservations are configured and managed through the cluster server software, and prevent other hosts from accessing particular volumes.

Unlike other types of reservations, a persistent reservation is used to perform the following functions:
* Reserve access across multiple host ports
* Provide various levels of access control
* Query the storage array about registered ports and reservations
* Provide for persistence of reservations in the event of a storage system power loss

The storage management software allows you to manage persistent reservations in the Subsystem Management window. The Persistent Reservation option enables you to perform the following tasks:
* View registration and reservation information for all volumes in the storage array
* Save detailed information about volume reservations and registrations
* Clear all registrations and reservations for a single volume or for all volumes in the storage array

For detailed procedures, see the Subsystem Management Window online help.

You can also manage persistent reservations through the script engine and the command line interface. For more information, see the Enterprise Management Window online help.

Configuring storage subsystem password protection

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

After you have set the password for each storage subsystem, you are prompted for that password the first time that 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. Configuring storage subsystems

This chapter 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 information, see the Subsystem Management window online help.

A logical drive is a logical structure that 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 either unconfigured capacity nodes or free capacity nodes in the storage subsystem from the Subsystem Management window. See Figure 11.
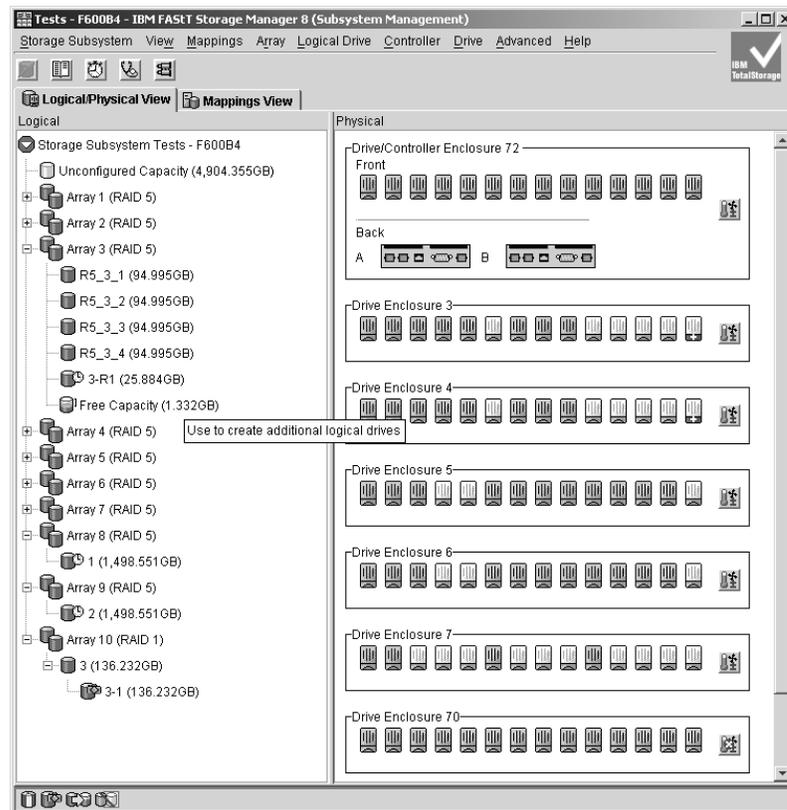
SJ001128

Figure 11. Unconfigured and free capacity nodes

If you have not configured any logical drives on the storage subsystem, the only node that is available is the unconfigured capacity node.

Creating logical drives from free or unconfigured capacity

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

1. Select either **Free Capacity Node** or **Unconfigured Capacity Node** and then click **Logical Drive Create**.

   The Create Logical Drive Wizard – Introduction window opens. See Figure 12.

   Note: If the Storage Partition feature is enabled, a Default Host Type window might open on top of the Create Logical Drive Wizard – Introduction window, that asks you to select a default host type for the logical drive or drives that you are going to be create. Select the appropriate host type from the pull-down list and click **OK** to close the Default Host Type window.
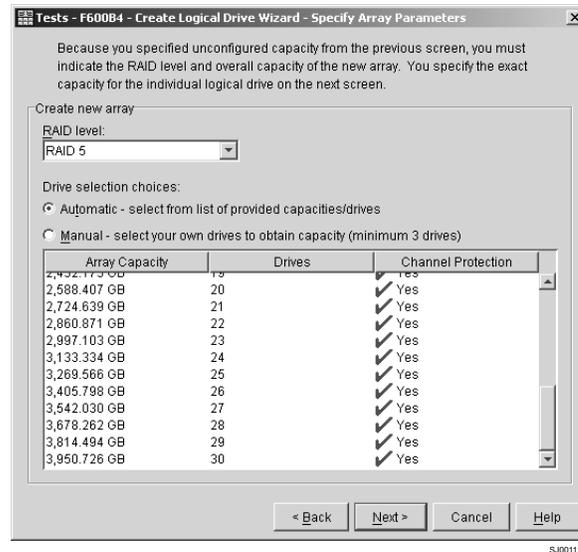
Figure 12. Create Logical Drive window

2. From the Create Logical Drive Wizard – Introduction window, select whether you want the new logical drive to be created from "free capacity" space in an existing defined array or arrays or from unconfigured optimal hard disk drives in the storage subsystem. Click **Next** to continue.

   Note: If there are multiple arrays that have free capacity in the storage subsystem, you can change the array by selecting free capacity before clicking **Next**.

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

3. If the new logical drive is going to be created from "free capacity" space in an existing defined array, click **Next** and go to step 4 on page 59. Otherwise, the Create Logical Drive Wizard – Specify Array Parameters window opens to allow you to define the parameters for an array that you will create for the logical drive. In this window, specify the following array parameters:

   - The RAID level: RAID-0, RAID-1, RAID-3 or RAID-5

   - The array capacity: Select the number of drives that will be used for this array. You can either specify the number of drives in the array and let the storage subsystem make the selection of which optimal hard drives in the storage subsystem to use, or you can manually select which hard drives will form the array.

Note: In manual mode, to select multiple drives, you can press and hold the **Ctrl** key and then use the mouse to select the desired hard drive or drives in the available unconfigured optimal hard drive list.

Click **Next** to continue. The Create Logical Drive Wizard - Specify Logical Drive Parameters window opens.

4. In the Create Logical Drive Wizard - Specify Logical Drive Parameters window, specify the following logical drive parameters for the logical drive you want to create:

Logical drive capacity
　　The new logical drive capacity can either be the entire unconfigured capacity in an array or a portion of the capacity.

Logical drive name
　　This name must be unique in the storage subsystem.

Settings
　　You can create the logical drive either using the FAStT Storage Server default settings or you can select your own customized I/O characteristics, controller ownership, and logical-drive-to-LUN mapping settings.

If you choose to use the FAStT Storage Server default settings, click **Finish** to create the logical drive, then go to step 6 on page 60. Otherwise, click **Next** to continue.

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

5. In the Create Logical Drive Wizard - Specify Advanced Logical Drive Parameters window, specify the following settings:

- The appropriate I/O characteristics (segment size and cache read-ahead multiplier)

  The I/O characteristics settings can be manually specified or automatically set based on the selection of one of the following expected logical drive usages: file system, database, or multimedia.

- The preferred controller ownership

- The logical-drive-to-LUN mapping setting

  The logical drive-to-LUN mapping preference can be one of the following two settings:

  Automatic
  　　The Automatic 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 (those that were designated by the default host group node in the Topology view). If the Storage Partition feature is not enabled, you must specify the Automatic setting.

  Map later with Storage Partitioning
  　　This setting specifies that you are not going to assign a LUN to the logical drive during creation. This setting enables you to define a specific logical drive-to-LUN mapping and create storage partitions using the Mappings Defined option. When you enable storage partitioning, specify this setting.

Click **Finish** to create the logical drive. The Create Logical Drive Wizard – Creation Successful window opens and states that the logical drive was successfully created.

6. If you do not want to create another logical drive, go to step 7. Otherwise, indicate in the Create Logical Drive Wizard – Creation Successful window if you want to create the logical drive from the same array, if there is free capacity, or from a different array if there is no capacity in the array in which the logical drive was just created. Click **YES** and go to step 4 on page 59.

7. If you do not want to create another logical drive, click **NO** and the Create Logical Drive Wizard – Completed window opens. Click **OK**.

8. 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 for your operating system to enable the operating system to discover the new logical drive.

Using available channel protection

When you create logical drives from unconfigured capacity, array candidates are shown in the Create Logical Drive pull-down menu. 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-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.

Storage partitioning

You can use the Storage Partitions feature of the 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 17 describes the storage partitioning terminology that is used in the Mappings window.

Table 17. Storage partitioning terminology

Term	Description
Storage partition	Storage partitions are 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.

Table 17. Storage partitioning terminology (continued)

Term	Description
Host port	Host ports physically reside on the host adapters and are automatically discovered by the Storage Manager software. To give a host computer access to a partition, you must define its associated host ports. The host ports request data from a logical drive on behalf of the host computer; therefore, without associated host ports, a host computer cannot be given a logical drive-to-LUN mapping or request data from a logical drive using a LUN. Initially, all discovered host ports belong to the default host group.
Host computer	A system that is directly attached to the storage subsystem through a fibre-channel I/O path. This system is used to serve data (typically in the form of files) from the storage subsystem. A system can be both a storage management station and a host simultaneously. **Note:** A defined host computer corresponds to a single computer that is 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.
Host group	A host group is 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. **Note:** You can define a host group that corresponds to a cluster or to a 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.
Default host group	A default host group is a logical collection of discovered host ports, defined host computers, and defined host groups in the storage-partition topology that fulfill the following requirements: • Are not involved in specific logical drive-to-LUN mappings • Share access to logical drives with default logical drive-to-LUN mappings
LUN	A LUN, logical unit number, is the number that a host computer uses to access a logical drive. Each host computer has its own LUN address space.
Specific logical drive-to-LUN mapping	The 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.

Table 17. Storage partitioning terminology (continued)

Term	Description
Default logical drive-to-LUN mapping	The default logical drive-to-LUN mapping enables host groups or host computers that do not have specific logical drive-to-LUN mappings (such as host computers or host groups that belong to the default host group) to access a particular logical drive. Logical drives are given default logical drive-to-LUN mappings when legacy logical drives (that are created using previous versions of the Storage Manager software) are automatically given default logical drive-to-LUN mappings. After installing Storage Manager Version 8.4 software, specific logical drive-to-LUN mappings are created.
Storage partitions mapping preference	You can choose one of the following storage partition mapping preferences when creating a logical drive: • Default logical drive-to-LUN mapping • No mapping Choose this option when you create storage partitions to define a specific logical drive-to-LUN mapping for this logical drive.

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 65 for more information. For detailed information about storage partitions, see the Subsystem Management window online help.
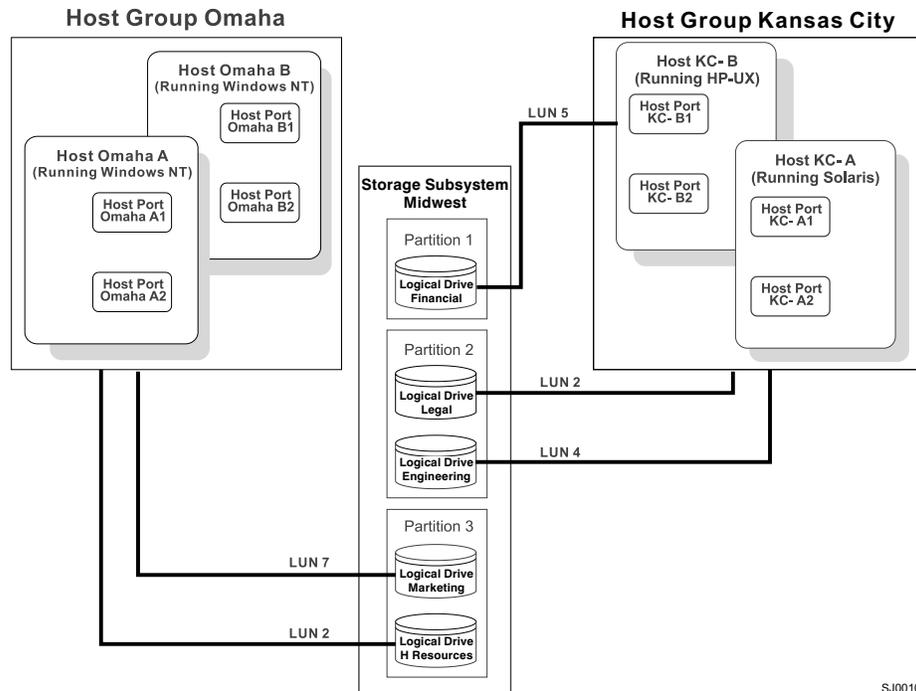
Figure 13. Storage partitioning example

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

In the example, 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'. The logical drive is accessed by host computers 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.

Obtaining a feature key

Depending on your FAStT storage server model, the storage partitioning feature might be enabled by default. If it is not enabled, you might have to contact IBM to purchase an option. 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.

Mappings window overview

Before you use storage partitions, you must define topological elements. To define topological elements, click the **Mappings** tab in the Subsystem Management window to display the Mappings window. Use the Mappings view 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, shown in Figure 14 on page 65, and Mappings. Table 18 describes the two Mappings views.

Table 18. Mappings window views

View	Description
Topology	The Topology view shows the defined topological elements (host ports, host computers, host groups) and the default host group. To use storage partitions you must define topological elements. If you are not using storage partitions, you can use the Topology view to define host computers.
Mappings	The Mappings view shows the default and specific logical drive-to-LUN mappings in a storage subsystem. The following information is displayed about the logical drives: • The topological entity that can access the logical drive • A logical drive name • A unique logical drive identifier • The LUN that is used to access the logical drive

Defining the storage partition topology

Click **Configure** → **Storage Partitioning** 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 65 shows the Topology view, which is located in the left

pane of the Mappings window.
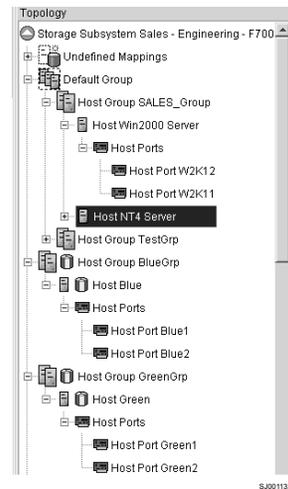
Figure 14. Topology view

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 'SALES_group' and 'TestGrp' and their associated host computers and host ports are not given specific logical drive-to-LUN mappings, so they are 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 are not part of 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. Click **Define New Mapping**.

    The Define New Mapping window opens, as shown in Figure 15 on page 66.
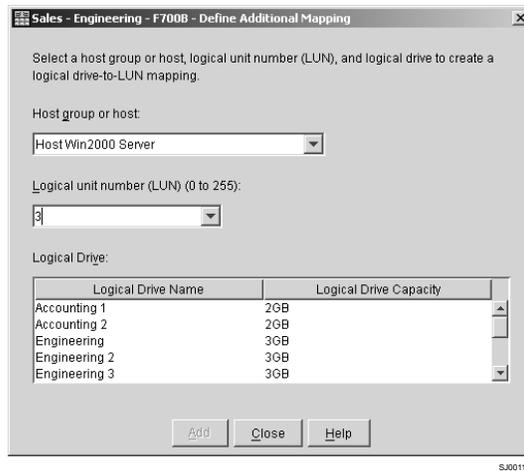
Figure 15. Define New Mapping window

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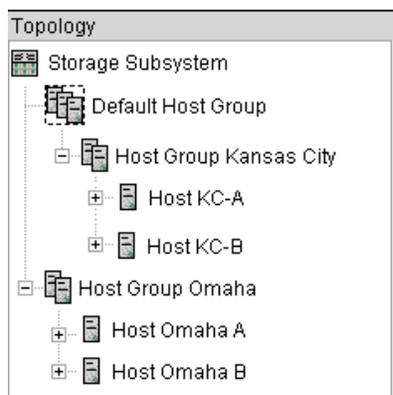
Figure 16. Creating a logical drive-to-LUN mapping in the Topology view

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

Heterogeneous hosts overview

The Heterogeneous Hosts feature enables host computers that are running different operating systems to access a single storage subsystem. Storage Manager firmware versions 04.00.xx.xx and earlier allowed only host computers that were running the same operating system to access a single storage subsystem.

Host computers can run 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 that is sending the information. For detailed information about defining heterogeneous host computer types, see the Subsystem Management window online help.

In a heterogeneous environment, you must set each host computer type to the appropriate operating system during host-port definition. This enables the firmware on each controller to respond correctly to the operating system of that host computer. See Figure 17 for an example.
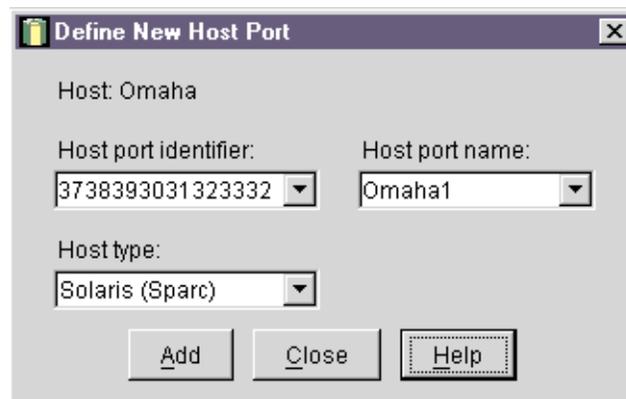
Figure 17. Define New Host Port window

Important: The Heterogeneous Hosts feature is available only when the Storage Partition feature is enabled.

In the example shown in Figure 18 on page 68, 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'.
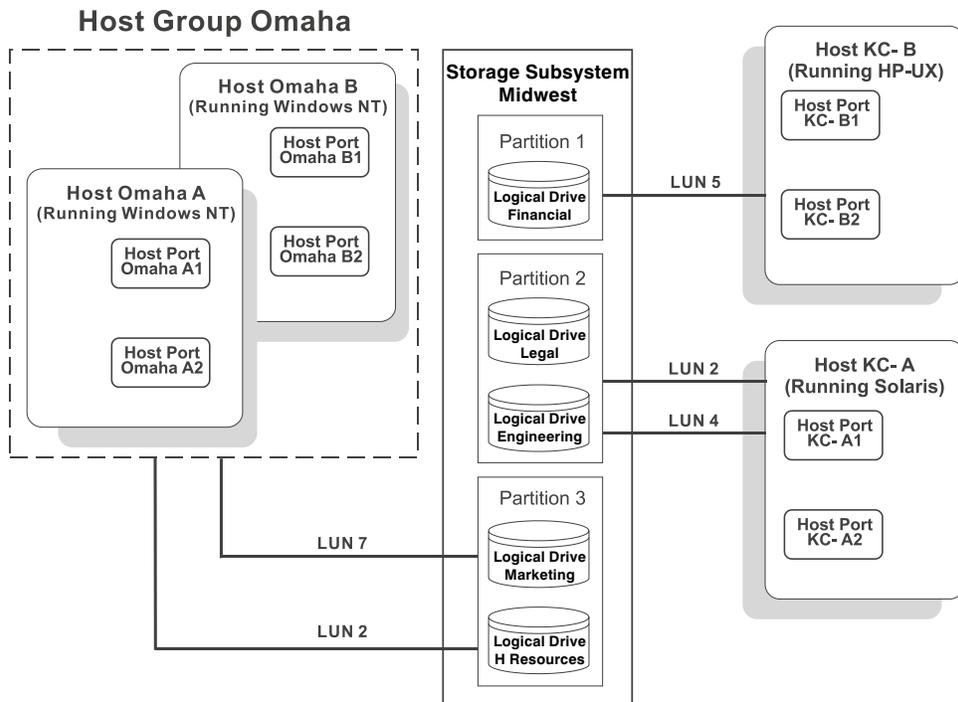
Figure 18. Heterogeneous hosts example

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 methods to display the host computer type of the host port in the Topology view:

- Place your cursor over the specific host port and the tooltip displays indicating 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 that are required for you to maintain and monitor 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 77.
- Recover from storage subsystem problems. See "Recovery Guru" on page 80.

Maintaining storage subsystems in a management domain

Use the Enterprise Management window to view the storage subsystem status icons and monitor the health of the storage subsystem. See Figure 19.
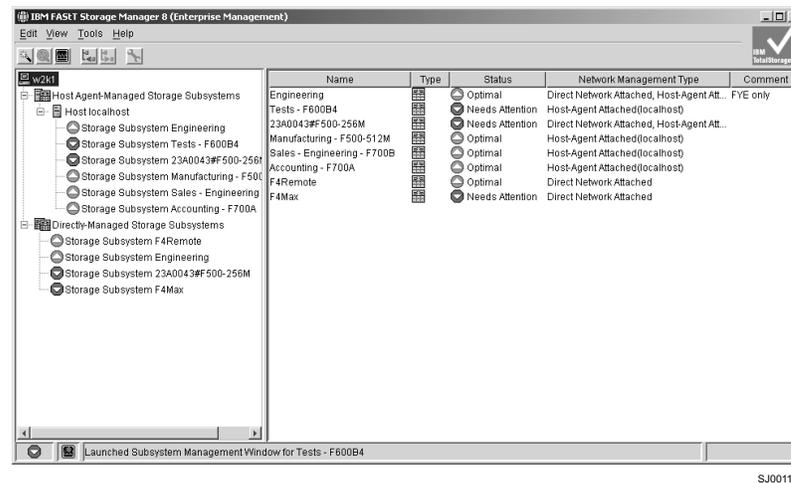
Figure 19. Monitoring storage subsystem health using the Enterprise Management window

Storage subsystem status quick reference

Table 19 on page 70 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 19. Storage subsystem status icon quick reference

Icon	Status	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 it. 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 to correct it.
🔧	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. **Note:** 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 controller or controllers in the storage subsystem over its network management connection. **Note:** 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. **Note:** 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 that the storage subsystem failed. The following list describes the various indicators:

- 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 in the Enterprise Management window
- The **Recovery Guru** 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 20 shows the Subsystem Management window Logical View and the failure icons for storage subsystem 'Engineering'.
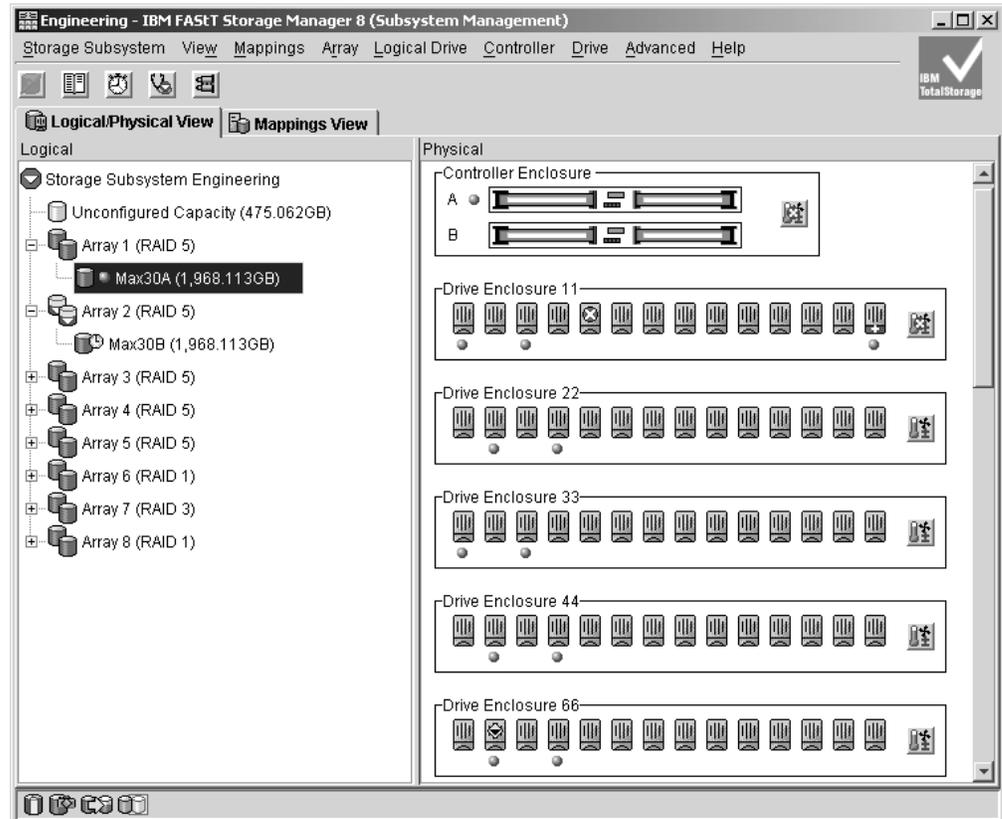
SJ001135

Figure 20. 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.

Updating the firmware in the storage subsystem and expansion units

The drive firmware and ESM firmware, like the controller firmware, need to be updated as required. The following section includes instructions for updating the firmware in the storage subsystem and expansion units.

Note: For Storage Manger 8.3 and later, you can perform ESM and drive firmware download using the Advanced menu in the Storage Subsystem window of the SMclient. For Storage Manager versions earlier than 8.3, you must use a separate Storage Manager Field tool program to perform the firmware downloads.

Downloading firmware

You must download firmware version 05.4x.xx before you download NVSRAM. You must have management connections to both controllers, and both controllers must be in an optimal state before you start the controller firmware and NVSRAM download.

Perform the following steps to download firmware version 05.4x.xx:

1. From the Enterprise Management window, select a storage subsystem.
2. Click **Tools** → **Manage Device**. The Subsystem Management window opens.
3. Click **Storage Subsystem** → **Download** → **Firmware**. Follow the instructions that display.
4. Continue with "Downloading NVSRAM."

Downloading NVSRAM

Perform the following steps to download NVSRAM:

1. From the Enterprise Management window, select a storage subsystem.
2. Click **Tools** → **Manage Device**. The Subsystem Management window opens.
3. Click **Storage Subsystem** → **Download** → **NVSRAM**. Follow the instructions that display.

Downloading drive firmware

Some of the firmware download can be done with I/O operations and some, like drive firmware, must be done with I/O quiesce. Use this option to transfer a downloadable fibre-channel hard drive firmware file to a drive or drives in the storage subsystem.

Stop all I/O operations to the storage subsystem before you begin to prevent application errors.

Attention: Note the following considerations before you download the drive firmware:

- The drive firmware files for various fibre-channel hard drive types are not compatible with each other. Ensure that the firmware that you download to the drives is compatible with the drives that you select. If noncompatible firmware is downloaded, the selected drives might become unusable, which will cause the logical drive to be in a degraded or even failed state.
- The drive firmware update must be performed without making any host I/O operations to the logical drives that are defined in the storage subsystem. Otherwise, it could cause the firmware download to fail and make the drive unusable, which could lead to loss of data availability.
- Do not make any configuration changes to the storage subsystem while downloading drive firmware or it could cause the firmware download to fail and make the selected drives unusable.
- If you download the drive firmware incorrectly, it could result in damage to the drives or loss of data.

Note: Downloads can take several minutes to complete. During a download, a Working dialog is displayed. Do not attempt another operation until the Working dialog is no longer displayed.

Perform the following steps to download the drive firmware:

1. From the Subsystem Management Window, select the **Advanced** → **Download** → **Drive Firmware**.

   The Download Drive Firmware window displays. The **Select Drives** field lists the details of all non-failed or optimal drives that are installed in the storage subsystem. The **Select File** field allows you to select a drive firmware file for download.

2. In the **Select Drives** field, highlight each drive that you want to download firmware to or click **Select All** to highlight all drives in the storage subsystem.

   Note: If you download firmware to more than one drive, all the drives that are selected must be from the same manufacturer and with the same product ID information.

3. Enter the firmware file to download in the **Select File** field by either typing the location and name of the file in the **Select File** text box (for example `C:\filename`), or clicking **Browse** and selecting the firmware file from a local or network drive.

4. Click **Start**.

   Note: The **Start** button will be unavailable until both a drive and a firmware file are selected.

5. Confirm your selections in the dialog and then click **Yes** to continue with the firmware download.

   The status in the **Select Drives** field changes from `Pending` to `Downloading` for the drive firmware operation in progress.

   The **Start** button changes to **Stop** after a firmware download is started. If you click **Stop** while a firmware download is in progress, the current download will finish before the operation stops. The **Status** field for the remaining drives changes to `Canceled`.

6. Monitor the progress and completion status of the download to the drives. The progress and status of each drive that is participating in the download is displayed in the **Status** field of the **Select Drives** field.

Note: If the file transfer fails, the status of the drive shows `Failed` and the rest of the selected drives that do not have new firmware downloaded yet shows `Canceled`.

Make sure that the new drive firmware file is compatible before you attempt another drive firmware download.

Environmental services monitor card

An environmental services monitor (ESM) card is a customer replaceable unit (CRU) component in the IBM FAStT storage expansion units (EXP700 and EXP500) that monitors the environmental condition of the components in that enclosure. It also provides the interface between the fibre-channel drives in a given expansion enclosure with other ESM cards and controller blades in a drive loop.

Downloading ESM firmware

With Storage Manager 8.4 and controller firmware 05.4x.xx.xx or higher, it is possible to update the ESM firmware while host I/O operations are made to the logical drives that are defined in the storage subsystem. The ESM firmware version must be the same in all of the fibre-channel drive enclosures of the same type in a given FAStT storage subsystem configuration.

For example, if the FAStT storage subsystem has three EXP700 fibre-channel drive enclosures and two EXP500 fibre-channel drive enclosures, the firmware of all the ESMs in the two EXP500 fibre-channel drive enclosures must be the same and the firmware of all ESMs in the three EXP700 fibre-channel drive enclosures must be the same.

The ESM code for the EXP500 fibre-channel drive enclosure is not compatible with the EXP700 fibre-channel drive enclosures.

Before you begin to download the ESM firmware, consider the following points:

- The IBM fibre-channel drive enclosures must be connected together in an IBM supported drive enclosure fibre-channel connection scheme.
- Both of the ESMs in each of the drive enclosures must be connected in dual redundant drive loops.
- Use SMclient to check for any loss of redundancy errors in the drive loop and to make the appropriate corrections before you attempt to download the ESM firmware.

Attention:

- Do not make any configuration changes to the storage subsystem while you download the ESM card firmware or it could cause the firmware download to fail and make the selected ESM cards unusable.
- If you download the ESM card firmware incorrectly, it could result in damage to the storage subsystem or loss of data.

Perform the following steps to download the ESM card firmware:

1. From the Subsystem Management window, select **Advanced → Download → Environmental (ESM) Card Firmware**.

   The Download Environmental (ESM) Card Firmware window is displayed. The **Select Enclosures** field lists the details of all the enclosures in the storage subsystem that contain ESM cards.The **Select File** field allows you to select an ESM card firmware file for download.

   Note: If an ESM card does not display in the list (because of a loss of redundancy or some other problem) run the Recovery Guru to diagnose and fix the problem before you continue with the download.

2. In the **Select Enclosures** field, highlight each enclosure to which you want to download firmware or click **Select All** to highlight all drive enclosures in the storage subsystem.

   Each drive enclosure that you select should have the same product ID.

   Important: If the storage subsystem has a combination of EXP500 and EXP700 fibre-channel drive enclosures, be careful when you select the drive enclosures that will be downloaded with the new firmware. The ESM code for the EXP500 fibre-channel drive enclosure is not compatible with the EXP700 fibre-channel drive enclosures.

3. Enter the firmware file to download in the **Select File** field by either typing the location and name of the file in the **Select File** text box (for example `C:\filename`), or clicking **Browse** and getting the firmware file from a local or network drive.

   Note: The **Browse** button will be unavailable until an enclosure has been selected.

4. Click **Start**.

   Note: The **Start** button will be unavailable until both a drive enclosure and a firmware file are selected.

5. Confirm your selections and then click **Yes** to continue with the firmware download.

The status in the **Select Enclosures** field changes from `Pending` to `Downloading` for the ESM card firmware operation in progress.

Note: If you click **Stop** while a firmware download is in progress, the current download will finish before the operation stops. The **Status** field for the remaining enclosures changes to `Canceled`.

6. Monitor the progress and completion status of the download to the enclosures. The progress and status of each drive enclosure that is participating in the download is displayed in the **Status** field of the **Select Enclosures** field.

   Note: Each firmware download can take several minutes to complete.

7. Click **Cancel** to close the Download Environmental (ESM) Card Firmware window when the ESM download status of all the selected drive enclosure changes to `Completed`. The ESM cards are now operating with the new firmware.

Note: If the file transfer fails, the status of one drive enclosure will show `Failed` and the rest of the selected drive enclosures that do not have new firmware downloaded yet will show `Canceled`.

Make sure that the new firmware file is compatible before you attempt another firmware download.

Viewing and recovering missing logical drives

A missing logical drive is a placeholder node that is displayed in the Logical view. It indicates that the storage subsystem has detected inaccessible drives that are associated with a logical drive. Typically, this results when you remove drives that are associated with an array, or when one or more drive enclosures lose power.

Missing logical drives are only displayed in the Logical view if they are standard logical drives or repository logical drives. In addition, one of the following conditions must exist:

* The logical drive has an existing logical drive-to-LUN mapping, and drives that are associated with the logical drive are no longer accessible.
* The logical drive is participating in a remote mirror as either a primary logical drive or a secondary logical drive, and drives that are associated with the logical drive are no longer accessible.
* The logical drive is a mirror repository logical drive, and drives that are associated with the logical drive are no longer accessible. The Recovery Guru has a special recovery procedure for this case. Two mirror repository logical drives are created together on the same array when the Remote Mirror Option feature is activated and one is used for each controller in the storage subsystem. If drives that are associated with the array are no longer accessible, then both mirror repository logical drives are missing, and all remote mirrors are in an Unsynchronized state.
* The logical drive is a base logical drive with associated FlashCopy logical drives, and drives that are associated with the logical drive are no longer accessible.
* The logical drive is a FlashCopy repository logical drive, and drives that are associated with the logical drive are no longer accessible.

If missing logical drives are detected by the storage subsystem, a Missing Logical Drives group is created in the Logical view of the Subsystem Management window.

Each missing logical drive is shown and identified by its worldwide name and logical drive type. Missing logical drives are identified as being one of the following types of drives:

- A standard logical drive
- A base logical drive
- A FlashCopy repository logical drive
- A primary logical drive
- A secondary logical drive
- A mirror repository logical drive

Missing logical drives, in most cases, are recoverable. Do not delete missing logical drives without confirming that the logical drives are no longer needed, because they will be permanently removed from the configuration.

If the storage subsystem detects that logical drives are missing because they have either been accidentally removed or their drive enclosures have sustained a power loss, you can recover these logical drives by using either of the following methods:

- Reinsert the drives back into the drive enclosure.
- Ensure that the power supplies of the drive enclosure are properly connected to an operating power source and have an optimal status.

Alert notification overview

This section provides a basis for understanding how alert notifications are sent. For detailed procedures, see 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 display 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 receive notifications for all storage subsystems that are 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 the fully qualified e-mail addresses (name@company.com). Enter all the addresses to which you want the information sent.
- On the SNMP traps page, type the community name and trap destination. The community name is set in the network management station (NMS) by the network administrator. The trap destination is the IP address or the NMS.

Important: To set up alert notifications using SNMP traps, you must copy and compile a management information base (MIB) file on the designated NMS. See 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 to you to configure alert notification destinations. You can set up alert-notification destination addresses to be notified 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 Storage Manager did not have the Event Monitor and the Enterprise Management window must have remained open to monitor the storage subsystems and receive alerts.

The Event Monitor runs continuously in the background. It monitors 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.

The Event Monitor is a separate program that is bundled with the Storage Manager client software.

Note: The Event Monitor cannot be installed without the client.

Install the Event Monitor on a management station or host computer that is 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 and SMclient send alerts to a remote system. The emwdata.bin file on the management station contains the name of the storage subsystem that is being monitored and the address where to send alerts. The alerts and errors that occur on the storage subsystem are continually monitored by SMclient and the Event Monitor. The Event Monitor takes over for the client after SMclient is shut down. When an event is detected, a notification is sent to the remote system.

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 SMclient and the Event Monitor software together. After the software is installed, the **Event Monitor** icon (shown in Figure 21 on page 79) 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 Destinations** 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.

If the Event Monitor is configured and running on more than one host computer or management station that is 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 Enterprise Management window share the information to send alert messages. The Enterprise Management window displays alert status to help you install and synchronize the Event Monitor. The parts of the Enterprise Management window that are related to event monitoring are shown in Figure 21 on page 79.

Synchronization button

Alert notification check mark
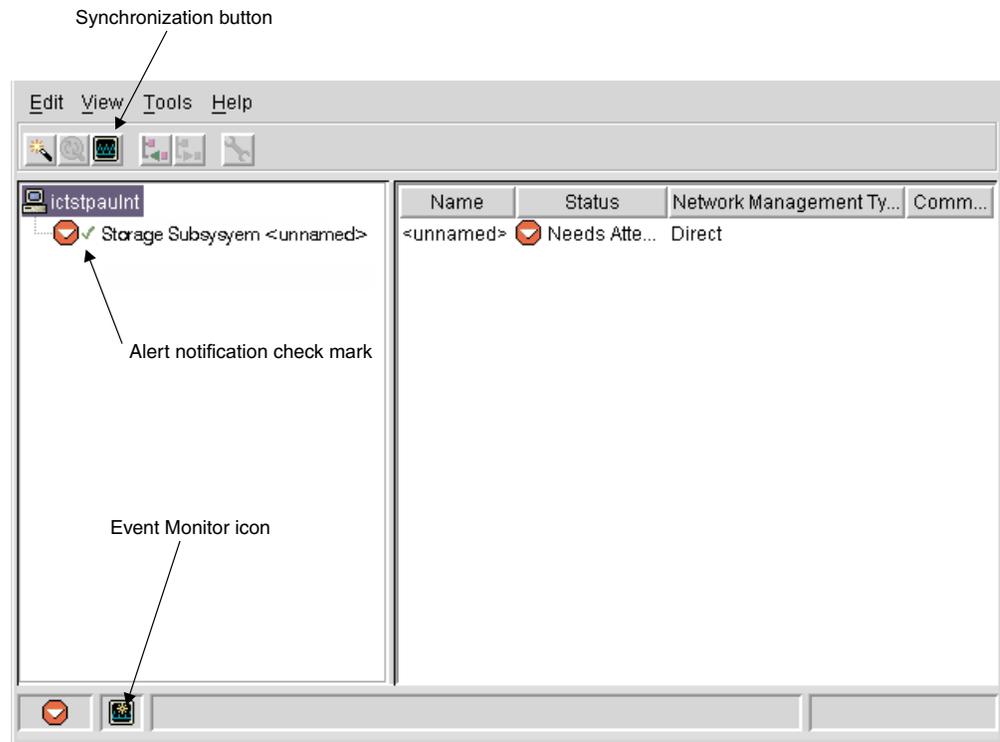
Event Monitor icon

Figure 21. Event monitoring example

Synchronizing the Enterprise Management window and Event Monitor

After the Event Monitor is installed, it continues to monitor storage subsystems and to 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, as shown in Figure 21.

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 **Synchronization** button synchronizes the Event Monitor and the Enterprise Management software components.

Note: The Enterprise Management window and the Event Monitor are automatically synchronized whenever you close 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

The Recovery Guru is a component of the Subsystem Management window in the SMclient package. The Recovery Guru diagnoses storage subsystem problems and suggests recovery procedures to correct the problems. To start the Recovery Guru, click **Recovery Guru** in the Subsystem Management window, shown in Figure 22, or click **Storage Subsystem → Recovery Guru**.
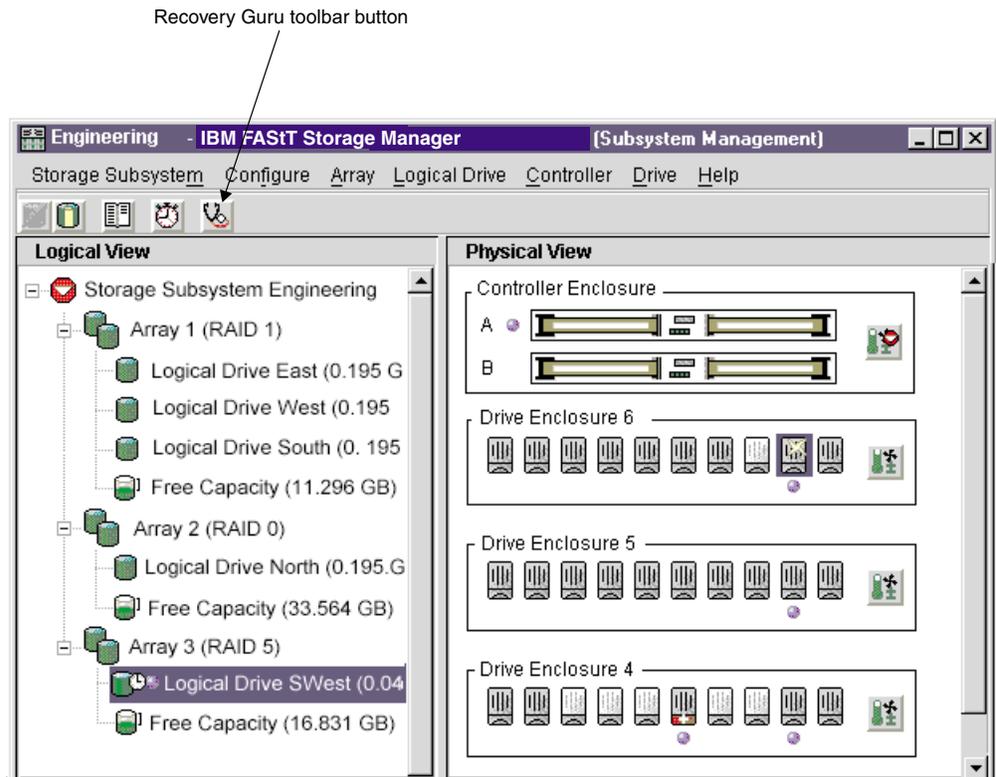
Figure 22. Location of the Recovery Guru toolbar button

The Recovery Guru window is shown in Figure 23 on page 81. The **Summary** pane shows that there are two different failures in this storage subsystem: a hot spare in use, and a failed battery CRU.

When you select a failure from the list in the **Summary** pane, the appropriate details and a recovery procedure display in the **Details** pane. For example, the Recovery Guru window 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 about this failure and how to recover from it.
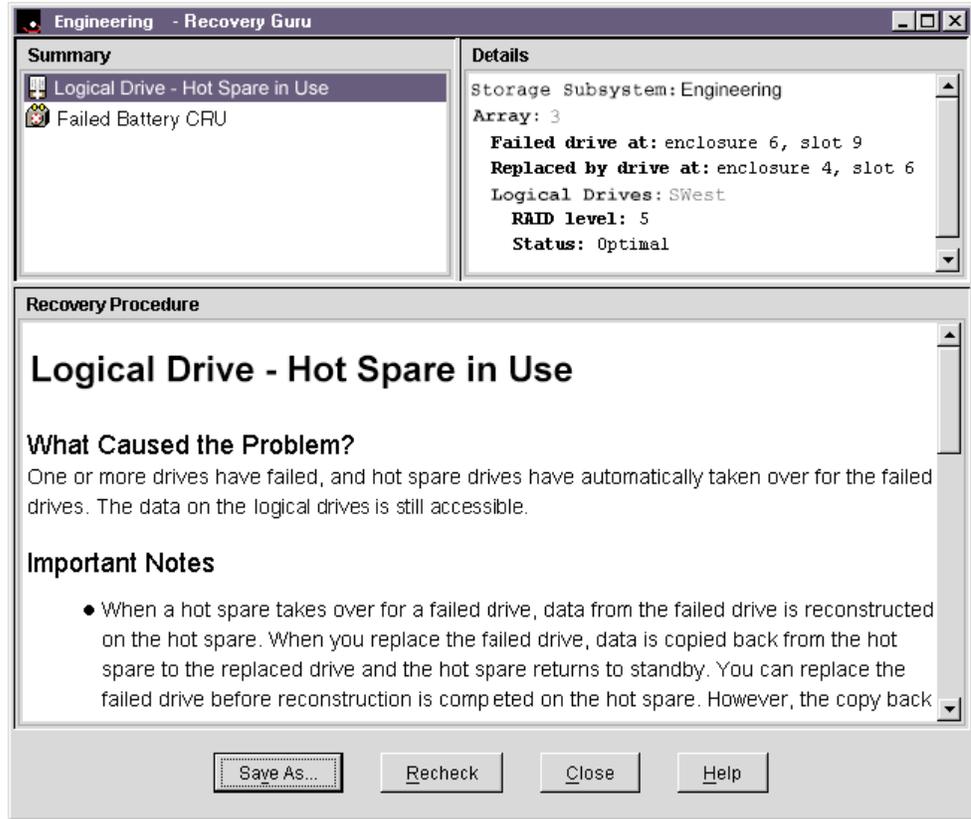
Figure 23. Recovery Guru window

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**. The data that is reconstructed to the hot-spare drive is copied back to the replaced physical drive. During the copyback operation, the status icon changes to **Replaced**, as shown in Figure 24 on page 82.

The drive icon changes from Failed to Replaced status.
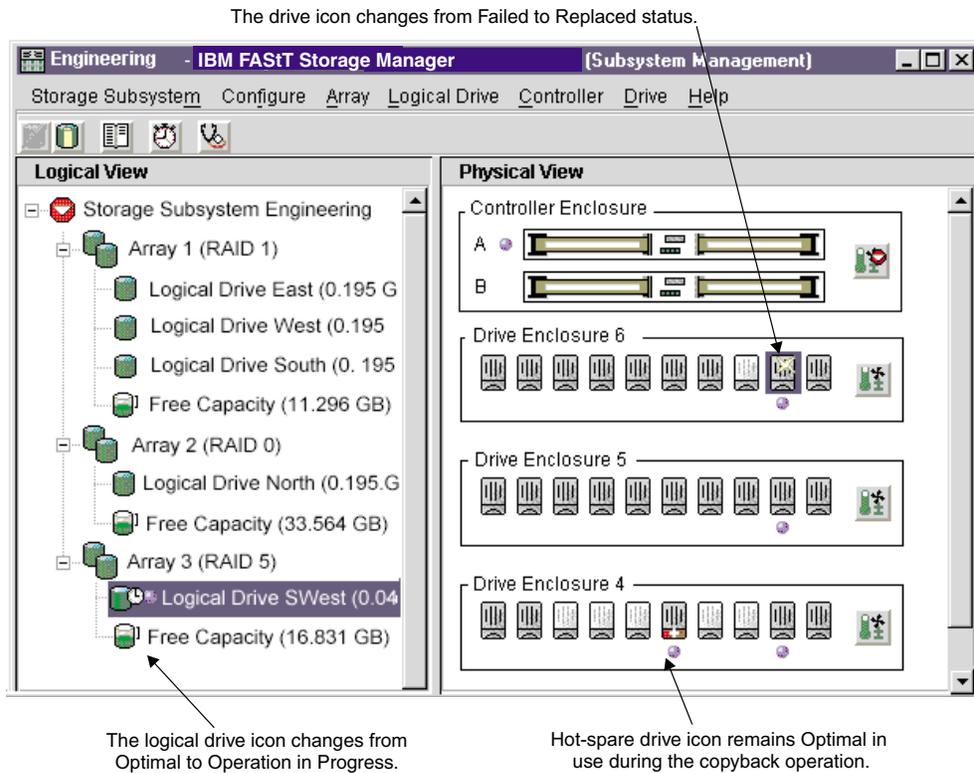
Figure 24. Recovery Guru window showing Replaced status icon

When the copyback operation is complete, the status icon changes to **Optimal**, as shown in Figure 25 on page 83.

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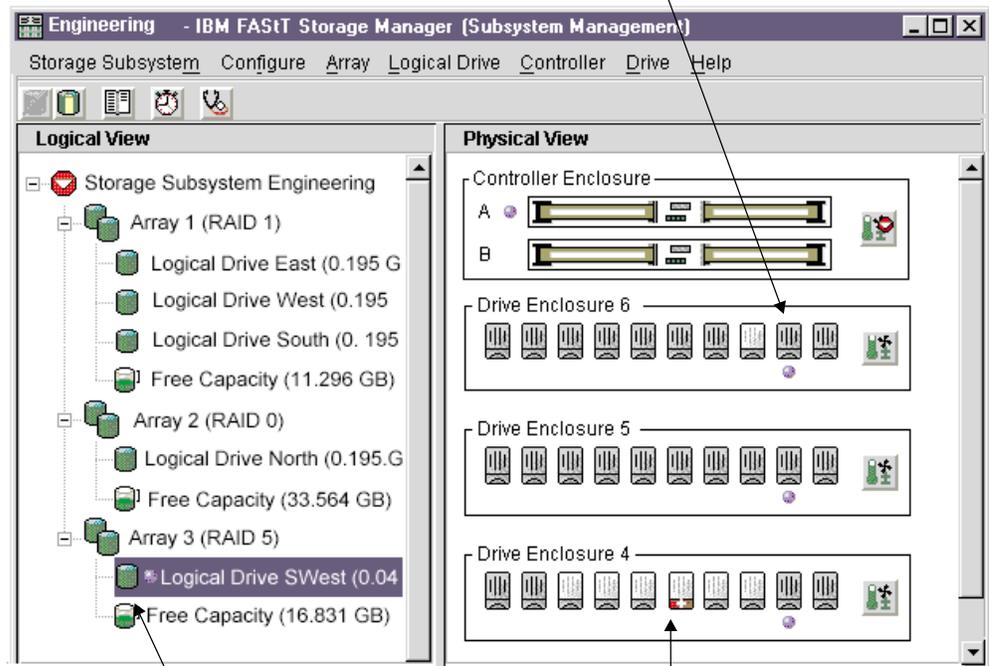
The logical drive icon changes from
Operation in Progress to Optimal.

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

Figure 25. Recovery Guru window showing Optimal status

After you correct the storage subsystem errors:

- The **Components** icon 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

The information in the chapter helps you use data from the Performance Monitor. This chapter also describes the tuning options that are available in Storage Manager 8.4 for optimizing storage subsystem and application performance. Use the Subsystem Management window Performance Monitor to monitor storage subsystem performance in real time and to 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 20 describes the Performance Monitor data that is displayed for selected devices.

Table 20. 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 86.
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 86.
Current KB per second	During the polling interval, the *transfer rate* 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 86.
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 86.
Current I/O per second	The average number of I/O requests that are 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 86.
Maximum I/O per second	The maximum number of I/O requests that are 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 86.

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 to 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 the data that is accessed during the polling session depends 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 data about I/O size to understand the maximum transfer rates possible. Then, use the tuning options that are available in Storage Manager 8.4 to optimize the I/O request rate to reach the maximum possible transfer rate.

Optimizing the fibre-channel I/O request rate

The fibre-channel I/O request rate can be affected by the following factors:
- The fibre-channel I/O access pattern (random or sequential) and I/O size
- The status of write-caching (enabled or disabled)
- The cache-hit percentage
- The RAID level
- The logical-drive modification priority
- The segment size
- The number of logical drives in the arrays or storage subsystem
- The fragmentation of files

   Note: Fragmentation affects logical drives with sequential fibre-channel I/O access patterns, not random fibre-channel I/O access patterns.

Determining the fibre-channel I/O access pattern and I/O size

To determine if the fibre-channel I/O access has sequential characteristics, enable 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 that the fibre-channel I/O has a sequential pattern. For more information, see "Optimizing the cache-hit percentage" on page 87. 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 the 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, you can enable cache read-ahead 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 the application behavior. Applications with a high read percentage perform well using RAID-5 logical drives because of the outstanding read performance of the RAID-5 configuration.

Applications with a low read percentage (write-intensive) do not perform as well on RAID-5 logical drives because of the way that a controller writes data and redundancy data to the drives in a RAID-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-5 to RAID-1 for faster performance.

Choosing an optimal logical-drive modification priority setting

The modification priority defines how much processing time is allocated for logical-drive modification operations versus system performance. The higher the priority, the faster the logical-drive modification operations 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 Storage Manager 7.01 and 7.02, the segment size is expressed in the number of data blocks. The segment size in Storage Manager 8.4 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 that are needed to satisfy an I/O request. This technique helps even more if you have random I/O access patterns. If you use a single logical drive for a single request, it 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 (which is the segment size multiplied by the number of logical drives in the array that are 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 time that you access a drive to read or write a file, it results in the movement of the read/write heads. Verify that the files on the logical drive are defragmented. When the files are defragmented, the data blocks that make up the files are next to each other, preventing extra read/write head movement when retrieving files. Fragmented files decrease the performance of a logical drive with sequential I/O access patterns.

Chapter 6. Critical event problem solving

When a critical event occurs, it is logged in the Event Log. It is also sent to any e-mail and SNMP trap destinations that you have configured. The critical event type and the sense key/ASC/ASCQ data are both shown in the event log details.

Table 21 provides more information about events with a critical priority, as shown in the Subsystem Management window event log.

Table 21. Critical events

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1001 - Channel failed	6/3F/C3	**Description:** 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. **Action:** Start the Recovery Guru and click the Failed Drive SCSI Channel recovery procedure. Contact your IBM technical-support representative to complete this procedure.
Event 1010 - Impending drive failure (PFA) detected	6/5D/80	**Description:** A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive might fail within 24 hours. **Action:** Start the Recovery Guru and click the Impending Drive Failure recovery procedure. Follow the instructions to correct the failure.
Event 1015 - Incorrect mode parameters set on drive	6/3F/BD	**Description:** 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. This indicates that the Qerr bit is set incorrectly on the drive specified in the FRU field of the Request Sense data. **Action:** The controller has not failed yet. Contact your IBM technical-support representative for the instructions to recover from this critical event.
Event 1207 - Fibre-channel link errors - threshold exceeded	None	**Description:** 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, damaged or faulty fibre-channel cables, or poor cable connections between components on the loop. **Action:** In the main Subsystem Management window, click **Help → Recovery Procedures**. Click **Fibre-channel Link Errors Threshold Exceeded** for more information about recovering from this failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1208 - Data rate negotiation failed	None	**Description:** The controller cannot auto-negotiate the transfer link rates. The controller considers the link to be down until negotiation is attempted at controller start-of-day, or when a signal is detected after a loss of signal. **Action:** Start the Recovery Guru and click the Data Rate Negotiation Failed recovery procedure. Follow the instructions to correct the failure.
Event 150E - Controller loopback diagnostics failed	None	**Description:** 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 occurs only on certain controller models. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 150F - Channel miswire	None	**Description:** Two or more drive channels are connected to the same fibre-channel loop. This can cause the storage subsystem to behave unpredictably. **Action:** Start the Recovery Guru and click the Channel Miswire recovery procedure. Follow the instructions to correct the failure.
Event 1510 - ESM canister miswire	None	**Description:** Two ESM canisters in the same drive enclosure are connected to the same fibre-channel loop. A level of redundancy has been lost and the I/O performance for this drive enclosure is reduced. **Action:** Start the Recovery Guru and click the ESM Canister Miswire recovery procedure. Follow the instructions to correct the failure.
Event 200A - Data/parity mismatch detected on logical drive	None	**Description:** A media scan operation has detected inconsistencies between a portion of the data blocks on the logical drive and the associated parity blocks. User data in this portion of the logical drive might have been lost. **Action:** Select an application-specific tool (if available) to verify that the data is correct on the logical drive. If no such tool is available, or if problems with the user data are reported, the entire logical drive contents should be restored from the most recent backup, if it is critical data.
Event 202E - Read drive error during interrupted write	3/11/8A	**Description:** A media error has occurred on a read operation during interrupted write processing. **Action:** Start the Recovery Guru and click the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2109 - Controller cache not enabled - cache sizes do not match	6/A1/00	**Description:** The controller cannot enable mirroring if the alternate controller cache size of both controllers is not the same. Verify that the cache size for both controllers is the same. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 210C - Controller cache battery failed	6/0C/80	**Description:** The controller has detected that the battery is not physically present, is fully discharged, or has reached its expiration date. **Action:** Start the Recovery Guru and click the Failed Battery CRU recovery procedure. Follow the instructions to correct the failure.
Event 210E - Controller cache memory recovery failed after power cycle or reset	6/0C/81	**Description:** Recovery from a data-cache error was unsuccessful. User data might have been lost. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 211B - Batteries present but NVSRAM configured for no batteries	None	**Description:** A battery is present in the storage subsystem but the NVSRAM is set to not include batteries. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 2110 - Controller cache memory initialization failed	6/40/81	**Description:** The controller has detected the failure of an internal controller component (RAID buffer). The internal controller component failure might have been detected during operation or during an on-board diagnostic routine. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 2113 - Controller cache battery nearing expiration	6/3F/D9	**Description:** The cache battery expiration is within the specified number of weeks. **Action:** Start the Recovery Guru and click the Battery Nearing Expiration recovery procedure. Follow the instructions to correct the failure.
Event 2229 - Drive failed by controller	None	**Description:** The controller failed a drive because of a problem with the drive. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 222D - Drive manually failed	6/3F/87	**Description:** The drive was manually failed by a user. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2247 - Data lost on the logical drive during unrecovered interrupted write	6/3F/EB	**Description:** An error has occurred during interrupted write processing during the start-of-day routine, which caused the logical drive to go into a failed state. **Action:** Start the Recovery Guru and click the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.
Event 2248 - Drive failed - write failure	6/3F/80	**Description:** The drive failed during a write command. The drive is marked failed. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 2249 - Drive capacity less than minimum	6/3F/8B	**Description:** 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. **Action:** Replace the drive with a larger capacity drive.
Event 224A - Drive has wrong block size	6/3F/8C	**Description:** The drive block size does not match that of the other drives in the logical drive. The drive is marked failed. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 224B - Drive failed - initialization failure	6/3F/86	**Description:** The drive failed either from a **Format Unit** command or a Write operation (issued when a logical drive was initialized). The drive is marked failed. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 224D - Drive failed - no response at start of day	6/3F/85	**Description:** The drive failed a **Read Capacity** or **Read** command during the start-of-day routine. The controller is unable to read the configuration information that is stored on the drive. The drive is marked failed. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 224E - Drive failed - initialization/reconstruction failure	6/3F/82	**Description:** The previously-failed drive is marked failed because of one of the following reasons: • The drive failed a **Format Unit** command that was issued to it • The reconstruction on the drive failed because the controller was unable to restore it (for example, because of an error that occurred on another drive that was required for reconstruction) **Action:** Start the Recovery Guru and follow the instructions to correct the failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2250 - Logical drive failure	6/3F/E0	**Description:** 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 that is protected by one drive. **Action:** Start the Recovery Guru and click the Failed Logical Drive Failure recovery procedure. Follow the instructions to correct the failure.
Event 2251 - Drive failed - reconstruction failure	6/3F/8E	**Description:** A drive failed because of a reconstruction failure during the start-of-day routine. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 2252 - Drive marked offline during interrupted write	6/3F/98	**Description:** An error has occurred during interrupted write processing which caused the logical drive to be marked failed. Drives in the array that did not experience the read error go into the offline state and log this error. **Action:** Start the Recovery Guru and click the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.
Event 2254 - Redundancy (parity) and data mismatch is detected	6/8E/01	**Description:** The controller detected inconsistent redundancy (parity) or data during a parity verification. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 2255 - Logical drive definition incompatible with ALT mode - ALT disabled	6/91/3B	**Description:** Auto-LUN transfer (ALT) works only with arrays that have only one logical drive defined. Currently there are arrays on the storage subsystem that have more than one logical drive defined; therefore, ALT mode has been disabled. The controller operates in normal redundant controller mode, and if there is a problem, it transfers all logical drives on an array instead of transferring individual logical drives. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 2602 - Automatic controller firmware synchronization failed	02/04/81	**Description:** 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 the firmware might cause unexpected results. **Action:** Try the firmware download again. If the problem persists, contact your IBM technical-support representative.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2801 - Storage subsystem running on uninterruptible power supply battery	6/3F/C8	**Description:** The uninterruptible power supply has indicated that ac power is no longer present and the uninterruptible power supply 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. **Action:** Start the Recovery Guru and click the Lost AC Power recovery procedure. Follow the instructions to correct the failure.
Event 2803 - Uninterruptible power supply battery - two minutes to failure	6/3F/C9	**Description:** The uninterruptible power supply has indicated that its standby power supply is nearing depletion. **Action:** Take actions to stop I/O activity to the controller. Normally, the controller changes from a write-back caching mode to a write-through mode.
Event 2804 - Uninterruptible power supply battery failed	None	**Description:** The uninterruptible power supply battery has failed. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 2807 - Environmental service monitor failed	None	**Description:** An ESM has failed. **Action:** Start the Recovery Guru and click the Failed Environmental Service Monitor CRU recovery procedure. Follow the instructions to correct the failure.
Event 2808 - Enclosure ID not unique	6/98/01	**Description:** 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. **Action:** Start the Recovery Guru and click the Enclosure ID Conflict recovery procedure. Follow the instructions to correct the failure.
Event 280A - Controller enclosure component missing	6/3F/C7	**Description:** 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. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 280B - Controller enclosure component failed	6/3F/C7	**Description:** 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. **Action:** Select Recovery Guru and follow the instructions to correct the failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 280D - Drive enclosure component failed	6/3F/C7	**Description:** 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. **Action:** Start the Recovery Guru and follow the instructions to correct the failure.
Event 280E - Standby power supply not fully charged	6/3F/CA	**Description:** The uninterruptible power supply has indicated that its standby power supply is not at full capacity. **Action:** Check the uninterruptible power supply to make sure that the standby power source (battery) is in working condition.
Event 280F - Environmental service monitor - loss of communication	6/E0/20	**Description:** Communication has been lost to one of the dual ESM CRUs in a drive enclosure. The drive enclosure has only one I/O path available. **Action:** Start the Recovery Guru and click the Environmental Service Monitor CRU - Loss of Communication recovery procedure. Follow the instructions to correct the failure.
Event 2813 - Minihub CRU failed	6/3F/C7	**Description:** Communication with the minihub CRU has been lost. This might be the result of a minihub CRU failure, a controller failure, or a failure in an internal backplane communications board. If there is only one minihub failure, the storage subsystem is still operational, but a second minihub failure could result in the failure of the affected enclosures. **Action:** Start the Recovery Guru and click the Failed Minihub CRU recovery procedure. Follow the instructions to correct the failure.
Event 2815 - GBIC failed	None	**Description:** 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. **Action:** Start the Recovery Guru and click the Failed GBIC recovery procedure. Follow the instructions to correct the failure.
Event 2816 - Enclosure ID conflict - duplicate IDs across drive enclosures	6/98/01	**Description:** Two or more drive enclosures are using the same enclosure identification number. **Action:** Start the Recovery Guru and click the Enclosure ID Conflict recovery procedure. Follow the instructions to correct the failure.
Event 2818 - Enclosure ID mismatch - duplicate IDs in the same drive enclosure	6/98/02	**Description:** A drive enclosure in the storage subsystem contains ESMs with different enclosure identification numbers. **Action:** Start the Recovery Guru and click the Enclosure ID Mismatch recovery procedure. Follow the instructions to correct the failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 281B - Nominal temperature exceeded	6/98/03	**Description:** 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. **Action:** Start the Recovery Guru and click the Nominal Temperature Exceeded recovery procedure. Follow the instructions to correct the failure.
Event 281C- Maximum temperature exceeded	6/3F/C6	**Description:** 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. **Action:** Start the Recovery Guru and click the Maximum Temperature Exceeded recovery procedure. Follow the instructions to correct the failure.
Event 281D - Temperature sensor removed	6/98/03	**Description:** A fan CRU containing a temperature sensor has been removed from the storage subsystem. **Action:** Replace the CRU as soon as possible. Start the Recovery Guru and click the Failed or Removed Fan CRU recovery procedure. Follow the instructions to correct the failure.
Event 281E - Environmental service monitor firmware mismatch	6/98/03	**Description:** A drive enclosure in the storage subsystem contains ESMs with different versions of firmware. ESMs 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. **Action:** Start the Recovery Guru and click the Environmental Service Monitor Firmware Version Mismatch recovery procedure. Follow the instructions to correct the failure.
Event 2821 - Incompatible minihub	None	**Description:** An incompatible minihub canister has been detected in the controller enclosure. **Action:** Start the Recovery Guru and click the Incompatible Minihub Canister recovery procedure. Follow the instructions to correct the failure.
Event 2829 - Controller redundancy lost	6/E0/20	**Description:** Communication has been lost between the two controllers through one of the drive loops (channels). **Action:** Start the Recovery Guru and see if there are other loss of redundancy problems being reported. If there are other problems being reported, fix those first. If you continue to have redundancy problems being reported, contact the IBM technical-support representative.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 282B - Drive enclosure path redundancy lost	6/E0/20	**Description:** A drive enclosure with redundant drive loops (channels) has lost communication through one of its loops. The enclosure has only one loop available for I/O. Correct this failure as soon as possible. Although the storage subsystem is still operational, a level of path redundancy has been lost. If the remaining drive loop fails, all I/O to that enclosure fails. **Action:** Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.
Event 282D - Drive path redundancy lost	6/E0/20	**Description:** A communication path with a drive has been lost. Correct this failure as soon as possible. The drive is still operational, but a level of path redundancy has been lost. If the other port on the drive or any other component fails on the working channel, the drive fails. **Action:** Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.
Event 3019 - Logical drive ownership changed due to failover	None	**Description:** 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. **Action:** Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.
Event 4011 - Logical drive not on preferred path	None	**Description:** The controller listed in the Recovery Guru area cannot be accessed. Any logical drives that have this controller assigned as their preferred path will be moved to the non-preferred path (alternate controller). **Action:** Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.
Event 5005 - Place controller offline	None	**Description:** 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 **Controller ▸ Run Diagnostics** menu option.) Or the controller is manually placed Offline using the **Controller ▸ Place Offline** menu option. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 502F - Missing logical drive deleted	None	**Description:** The storage subsystem has detected that the drives that are associated with a logical drive are no longer accessible. This can be the result of removing all the drives that are associated with an array or a loss of power to one or more drive enclosures. **Action:** Start the Recovery Guru and click the Missing Logical Drive recovery procedure. Follow the instructions to correct the failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 5038 - Controller in lockout mode	None	**Description:** Both controllers have been placed in lockout mode for 10 minutes because password authentication failures have exceeded 10 attempts within a ten-minute period. During the lockout period, both controllers will deny all authentication requests. When the 10-minute lockout expires, the controller resets the total authentication failure counter and unlocks itself. **Action:** Wait 10 minutes and try to enter the password again.
Event 5602 - This controller's alternate failed - timeout waiting for results	None	**Description:** 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. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 560B - CtlrDiag task cannot obtain Mode Select lock	None	**Description:** This controller is attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 560C - CtlrDiag task on controller's alternate cannot obtain Mode	None	**Description:** 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. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 560D - Diagnostics read test failed on controller	None	**Description:** While running diagnostics, the controller detected that the information that was 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 that is being read. The controller is placed offline as a result of this failure. **Action:** Select the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 560E - This controller's alternate failed diagnostics read test	None	**Description:** 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 that is being read. The alternate controller in this pair is placed offline. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 560F - Diagnostics write test failed on controller	None	**Description:** 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 that is being written. The controller is placed offline. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 5610 - This controller's alternate failed diagnostics write test	None	**Description:** 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 that is being written. The alternate controller in this pair is placed offline. **Action:** Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.
Event 5616 - Diagnostics rejected - configuration error on controller	None	**Description:** 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. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 5617 - Diagnostics rejected - configuration error on controller's alternate	None	**Description:** 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. **Action:** Contact your IBM technical-support representative for the instructions to recover from this failure.
Event 6101 - Internal configuration database full	None	**Description:** Because of the amount of data that is required to store certain configuration data, the maximum number of logical drives has been underestimated. One or both of the following types of data might have caused the internal configuration database to become full: • FlashCopy logical drive configuration data • Remote Mirror configuration data **Action:** To recover from this event, you can delete one or more FlashCopy logical drives from your storage subsystem or you can remove one or more Remote Mirror relationships.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6200 - FlashCopy repository logical drive threshold exceeded	None	**Description:** The FlashCopy repository logical drive capacity has exceeded a warning threshold level. If the capacity of the FlashCopy repository logical drive becomes full, its associated FlashCopy logical drive can fail. This is the last warning that you receive before the FlashCopy repository logical drive becomes full. **Action:** Start the Recovery Guru and click the FlashCopy Repository Logical Drive Threshold Exceeded recovery procedure. Follow the instructions to correct this failure.
Event 6201 - FlashCopy repository logical drive full	None	**Description:** All of the available capacity on the FlashCopy repository logical drive has been used. The failure policy of the FlashCopy repository logical drive determines what happens when the FlashCopy repository logical drive becomes full. The failure policy can be set to either fail the FlashCopy logical drive (default setting) or fail incoming I/Os to the base logical drive. **Action:** Start the Recovery Guru and click the FlashCopy Repository Logical Drive Capacity - Full recovery procedure. Follow the instructions to correct this failure.
Event 6202 - Failed FlashCopy logical drive	None	**Description:** Either the FlashCopy repository logical drive that is associated with the FlashCopy logical drive is full or its associated base or FlashCopy repository logical drives have failed due to one or more drive failures on their respective arrays. **Action:** Start the Recovery Guru and click the Failed FlashCopy Logical Drive recovery procedure. Follow the instructions to correct this failure.
Event 6400 - Dual primary logical drive	None	**Description:** Both logical drives have been promoted to a primary logical drive after a forced role reversal. This event might be reported when the controller resets or when a cable from an array to a fibre-channel switch is reinserted after it was removed and the other logical drive was promoted to a primary logical drive. **Action:** Start the Recovery Guru and click the Dual Primary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.
Event 6401 - Dual secondary logical drive	None	**Description:** Both logical drives in the Remote Mirror have been demoted to secondary logical drives after a forced role reversal. This could be reported when the controller resets or when a cable from an array to a fibre-channel switch is reinserted after it was removed and the other logical drive was promoted to a secondary logical drive. **Action:** Start the Recovery Guru and click the Dual Secondary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.

Table 21. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6402 - Mirror data unsynchronized	Not recorded with event	**Description:** This might occur because of I/O errors but there should be other events associated with it. One of the other errors is the root cause, that contains the sense data. A **Needs Attention** icon displays on both the primary and secondary storage subsystems of the Remote Mirror. **Action:** Start the Recovery Guru and click the Mirror Data Unsynchronized recovery procedure. Follow the instructions to correct this failure.
Event 6503 - Remote logical drive link down	None	**Description:** This event is triggered when either a cable between one array and its peer has been disconnected, the fibre-channel switch has failed, or the peer array has reset. This error could result in the Mirror Data Unsynchronized, event 6402. The affected remote logical drive displays an **Unresponsive** icon, and this state will be selected in the tooltip when you pass your cursor over the logical drive. **Action:** Start the Recovery Guru and click the Mirror Communication Error - Unable to Contact Logical Drive recovery procedure. Follow the instructions to correct this failure.
Event 6505 - WWN change failed	None	**Description:** Mirroring causes a WWN change to be communicated between arrays. Failure of a WWN change is caused by non-I/O communication errors between one array, on which the WWN has changed, and a peer array. (The array WWN is the unique name that is used to locate an array on a fibre network. When both controllers in an array are replaced, the array WWN changes). The affected remote logical drive displays an **Unresponsive** icon and this state will be selected in the tooltip when you pass your cursor over the logical drive. **Action:** Start the Recovery Guru and click the Unable to Update Remote Mirror recovery procedure. Follow the instructions to correct this failure. The only solution to this problem is to delete the Remote Mirror and then to establish another one.
Event 6600 - Logical drive copy operation failed	None	**Description:** A logical drive copy operation with a status of In Progress has failed. This failure can be caused by a read error on the source logical drive, a write error on the target logical drive, or because of a failure that occurred on the storage subsystem that affects the source logical drive or target logical drive. **Action:** Start the Recovery Guru and click the Logical Drive Copy Operation Failed recovery procedure. Follow the instructions to correct this failure.

Appendix A. Online help 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.

Populating a management domain

See the following online help	For information about the following task
Enterprise Management window	Adding a device to a management domain
	Correcting a partially managed device
	Discovering a newly attached host-agent managed storage subsystem
	Performing an initial auto-discovery
	Recovering from damaged configuration files
	Removing a device from a management domain

Configuring storage subsystems

See the following online help	For information about the following task
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

Using the Script Editor

See 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
	Running the currently loaded script
	Interpreting script results
	Opening a new script
	Saving the script results to a local file
	Saving the script in the Script view
	Using the Script Editor
	Verifying the syntax of the currently loaded script

Configuring storage partitions

See 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
	Undefining a host port
	Viewing a list of discovered host ports that are not defined

Protecting data

See 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

Event notification

See 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 that is 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

Recovering from problems

See 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

Miscellaneous system administration

See the following online help	For information about the following task
Subsystem Management window	Listing logical or physical components that are 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

Security

See the following online help	For information about the following task
Subsystem Management window	Changing a storage subsystem password
	Entering a storage subsystem password

See the following online help	For information about the following task
Enterprise Management window	Using passwords in the Script Editor

Performance and tuning

See the following online help	For information about the following task
Subsystem Management window	Changing the 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
	Changing 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 the cache properties of a logical drive
	Specifying the storage subsystem cache settings

Appendix B. Additional FAStT documentation

The following tables present an overview of the FAStT900, FAStT700, and FAStT600 Fibre Channel Storage Server product libraries, as well as other related documents. Each table lists documents that are included in the libraries and what common tasks they address. Click on active links in the tables to access those documents currently available on the Internet. You can access documentation for the other FAStT products at the following Web site:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt/

FAStT900 Fibre Channel Storage Server library

Table 22 associates each document in the FAStT900 Fibre Channel Storage Server library with its related common user tasks.

Table 22. TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM TotalStorage FAStT900 Installation and Support Guide, GC26-7530	✔	✔		✔		
IBM TotalStorage FAStT900 Fibre Channel Cabling Instructions, 24P8135	✔	✔				
IBM TotalStorage FAStT900 Storage Server User's Guide, GC26-7534				✔	✔	✔
IBM TotalStorage FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✔			✔	
IBM FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✔			✔	
IBM TotalStorage FAStT900 Rack Mounting Instructions, 19K0900	✔	✔				

Table 22. TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks (continued)

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329	✔	✔			✔	✔
IBM FAStT Management Suite Java User's Guide, 32P0081					✔	✔
IBM TotalStorage FAStT Fibre Channel Hardware Maintenance Manual and Problem Determination Guide, GC26-7528						✔

FAStT700 Fibre Channel Storage Server library

Table 23 associates each document in the FAStT700 Fibre Channel Storage Server library with its related common user tasks.

Table 23. TotalStorage FAStT700 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM FAStT700 Installation and Support Guide, 32P0171	✔	✔		✔		
IBM FAStT700 Fibre Channel Cabling Instructions, 32P0343	✔	✔				
IBM FAStT700 Fibre Channel Storage Server User's Guide, 32P0341				✔	✔	✔
IBM EXP700 Storage Expansion Unit Installation and User's Guide, 32P0178	✔	✔		✔	✔	✔
IBM FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✔			✔	
IBM TotalStorage FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✔			✔	
IBM FAStT Management Suite Java User's Guide, 32P0081					✔	✔
IBM Fibre Channel Hardware Maintenance Manual, 19K6130						✔
IBM Fibre Channel Problem Determination Guide, 48P9804						✔

FAStT600 Fibre Channel Storage Server library

Table 24 associates each document in the FAStT600 Fibre Channel Storage Server library with its related common user tasks.

Table 24. TotalStorage FAStT600 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM TotalStorage FAStT600 Fibre Channel Storage Server Installation and User's Guide, GC26-7531	✔	✔		✔		
IBM TotalStorage FAStT Fibre Channel Hardware Maintenance Manual and Problem Determination Guide, GC26-7528						✔
IBM TotalStorage FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✔			✔	
IBM TotalStorage FAStT600 Rack Mounting Instructions, 24P8125	✔	✔				
IBM TotalStorage FAStT600 Cabling Instructions, 24P8126	✔	✔				

FAStT–related documents

Table 25 associates each of the following documents related to FAStT operations with its related common user tasks.

Table 25. TotalStorage FAStT related document titles by user tasks

Title	User Tasks					
	Planning	**Hardware Installation**	**Software Installation**	**Configuration**	**Operation and Administration**	**Diagnosis and Maintenance**
IBM Safety Information, P48P9741					✔	
IBM Netfinity® Fibre Channel Cabling Instructions, 19K0906		✔				
IBM Fibre Channel SAN Configuration Setup Guide, 25P2509	✔		✔	✔	✔	

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.

Trademarks

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

IBM
AIX
e-server logo
FlashCopy
HelpCenter
Intellistation

115

Netfinity
ServerProven
TotalStorage
xSeries

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.

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.

Glossary

This glossary provides definitions for the terminology used for the IBM TotalStorage FAStT. This glossary also provides definitions for the terminology used for the IBM TotalStorage FAStT Storage Manager.

This glossary defines technical terms and abbreviations used in this document. If you do not find the term you are looking for, see the *IBM Glossary of Computing Terms* located at www.ibm.com/networking/nsg/nsgmain.htm

This glossary also includes terms and definitions from:

- *Information Technology Vocabulary* by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- *IBM Glossary of Computing Terms*. New York: McGraw-Hill, 1994.

The following cross-reference conventions are used in this glossary:

See Refers you to (a) a term that is the expanded form of an abbreviation or acronym, or (b) a synonym or more preferred term.

See also
Refers you to a related term.

Abstract Windowing Toolkit (AWT). A Java graphical user interface (GUI).

accelerated graphics port (AGP). A bus specification that gives low-cost 3D graphics cards faster access to main memory on personal computers than the usual peripheral component interconnect (PCI) bus. AGP reduces the overall cost of creating high-end graphics subsystems by using existing system memory.

access volume. A special logical drive that allows the host-agent to communicate with the controllers in the storage subsystem.

adapter. A printed circuit assembly that transmits user data input/output (I/O) between the internal bus of the host system and the external fibre-channel (FC) link and vice versa. Also called an I/O adapter, host adapter, or FC adapter.

ADT. See *auto-logical drive transfer feature*.

advanced technology (AT) bus architecture. A bus standard for IBM compatibles. It extends the XT bus architecture to 16 bits and also allows for bus mastering, although only the first 16 MB of main memory are available for direct access.

agent. A server program that receives virtual connections from the network manager (the client program) in a Simple Network Management Protocol-Transmission Control Protocol/Internet Protocol (SNMP-TCP/IP) network-managing environment.

AGP. See *accelerated graphics port*.

AL_PA. See *arbitrated loop physical address*.

arbitrated loop. A shared 100 MBps fibre-channel transport structured as a loop and supporting up to 126 devices and one fabric attachment. A port must successfully arbitrate before a circuit can be established.

arbitrated loop physical address (AL_PA). One of three existing fibre-channel topologies, in which two to 126 ports are interconnected serially in a single loop circuit. Access to the fibre channel-Arbitrated Loop (FC-AL) is controlled by an arbitration scheme. The FC-AL topology supports all classes of service and guarantees in-order delivery of FC frames when the originator and responder are on the same FC-AL. The default topology for the disk array is arbitrated loop. An arbitrated loop is sometimes referred to as Stealth Mode.

auto-logical drive transfer (ADT) feature.. ADT is a built-in feature of the controller firmware that allows logical drive-level failover rather than controller-level failover.

auto volume transfer/auto disk transfer (AVT/ADT). A function that provides automatic failover in case of controller failure on a storage subsystem.

AVT/ADT. See *auto volume transfer/auto disk transfer*.

AWT. See *Abstract Windowing Toolkit*.

basic input/output system (BIOS). Code that controls basic hardware operations, such as interactions with diskette drives, hard disk drives, and the keyboard.

BIOS. See *basic input/output system*.

BOOTP. See *bootstrap protocol*.

bootstrap protocol (BOOTP). A Transmission Control Protocol/Internet Protocol (TCP/IP) protocol that a diskless workstation or network computer uses to obtain its Internet Protocol (IP) address and other network information such as server address and default gateway.

bridge. A storage area network (SAN) device that provides physical and transport conversion, such as fibre channel to small computer system interface (SCSI) bridge.

bridge group. A bridge and the collection of devices connected to it. Bridge Groups are discovered by the SANavigator tool and are displayed with a gray background on the Physical and Data Path Maps.

broadcast. A method of sending a Simple Network Management Protocol (SNMP) request for information to all the devices on a subnet that use a single special request. Because of its efficiency, the SANavigator tool sets its default method of discovery to broadcast. However, a network administrator might disable this method on the network router.

cathode ray tube (CRT). An electrical device for displaying images by exciting phosphor dots with a scanned electron beam. CRTs are found in computer video display units (VDUs) and monitors, televisions, and oscilloscopes.

CDPD. See *cellular digital packet data*.

cellular digital packet data (CDPD). A wireless standard that provides two-way, 19.2 kps packet data transmission over existing cellular telephone channels.

CGA. See *color graphics adapter*.

client. A computer system or process that requests a service of another computer system or process that is typically referred to as a server. Multiple clients can share access to a common server.

color graphics adapter (CGA). An early, now obsolete, IBM video display standard for use on IBM personal computers. CGA displays 80 x 25 or 40 x 25 text in 16 colors, 640 x 200 pixel graphics in two colors or 320 x 200 pixel graphics in four colors.

command. Any selection on a dialog box or elsewhere in the user interface that causes the SANavigator tool to perform a task.

community strings. The name of a community contained in each Simple Network Management Protocol (SNMP) message. SNMP has no standard mechanisms for verifying that a message was sent by a member of the community, keeping the contents of a message private, or for determining if a message has been changed or replayed.

CRC. See *cyclic redundancy check*.

CRT. See *cathode ray tube*.

cyclic redundancy check (CRC). (1) A redundancy check in which the check key is generated by a cyclic algorithm. (2) An error detection technique performed at both the sending and receiving stations.

dac. See *disk array controller*.

dar. See *disk array router*.

DASD. See *direct-access storage device*.

default host group. A logical collection of discovered host ports, defined host computers, and defined host groups in the storage-partition topology that fulfill the following requirements:

- Are not involved in specific logical drive-to-LUN mappings
- Share access to logical drives with default logical drive-to-LUN mappings

device type. Identifier used to place devices in the physical map, such as the switch, hub, storage.

DHCP. See *Dynamic Host Configuration Protocol*.

direct access storage device (DASD). IBM mainframe terminology for a data storage device by which information can be accessed directly, instead of bypassing sequentially through all storage areas. For example, a disk drive is a DASD, in contrast with a tape drive, which stores data as a linear sequence.

direct memory access (DMA). The transfer of data between memory and an input/output (I/O) device without processor intervention.

disk array controller (dac). A disk array controller device that represents the two controllers of an array. See also *disk array router*..

disk array router (dar). A disk array router that represents an entire array, including current and deferred paths to all logical unit numbers (LUNs) (hdisks on AIX). See also *disk array controller*.

DMA. See *direct memory access*.

domain. The most significant byte in the node port (N_port) identifier for the fibre-channel (FC) device. It is not used in the fibre channel-small computer system interface (FC-SCSI) hardware path ID. It is required to be the same for all SCSI targets logically connected to an FC adapter.

DRAM. See *dynamic random access memory*.

dynamic host configuration protocol (DHCP). A protocol defined by the Internet Engineering Task Force that is used for dynamically assigning IP addresses to computers in a network.

dynamic random access memory (DRAM). A storage in which the cells require repetitive application of control signals to retain stored data.

E_port. See *expansion port*.

ECC. See *error correction coding*.

EEPROM. See *electrically erasable programmable read-only memory*.

EGA. See *enhanced graphics adapter*.

electrically erasable programmable read-only memory (EEPROM). A type of non-volatile storage device that can be erased with an electrical signal. Writing to EEPROM takes much longer than reading. It also can only be reprogrammed a limited number of times before it wears out. Therefore, it is appropriate for storing small amounts of data that are changed infrequently.

electrostatic discharge (ESD). The flow of current that results when objects that have a static charge come into close enough proximity to discharge.

enhanced graphics adapter (EGA). An IBM video display standard that provides text and graphics with a resolution of 640 x 350 pixels of 16 colors. It emulates the Color Graphics Adapter (CGA) and the Monochrome Display Adapter (MDA) and was superseded by the Video Graphics Display (VGA).

enhanced small device interface (ESDI). A hard disk controller standard that allows disks to communicate with computers at high speeds. ESDI drives typically transfer data at about 10 megabits per second, although they are capable of doubling that speed.

error correction coding (ECC). A method for encoding data so that transmission errors can be detected and corrected by examination of the data on the receiving end. Most ECCs are characterized by the maximum number of errors they can detect and correct.

error detection coding. A method for encoding data so that errors that occur during storage or transmission can be detected. Most error detection codes are characterized by the maximum number of errors they can detect. The simplest form of error detection is by using a single added parity bit or a cyclic redundancy check. Adding multiple parity bits can detect not only that an error has occurred, but also which bits have been inverted, thereby indicating which bits should be re-inverted to restore the original data.

ESD. See *electrostatic discharge*.

ESDI. See *enhanced small device interface*.

EXP. See *expansion unit*.

expansion port (E_port). An expansion port (E_port) that connects the switches for two fabrics (also used for McData ES-1000 B ports).

expansion unit. A feature that can be connected to a system unit to provide additional storage and processing capacity.

eXtended graphics array (XGA). An IBM advanced standard for graphics controller and display mode design introduced in 1990. XGA, used mostly on workstation-level systems, supports a resolution of 1024 x 768 pixels with a palette of 256 colors, or 640 x 480 with high color (16 bits per pixel). XGA-2 added 1024 x 768 support for high color and higher refresh rates, improved performance, and supports 1360 x 1024 in 16 colors.

F_port. See *fabric port*.

fabric group. A collection of interconnected storage area network (SAN) devices discovered by the SANavigator tool and displayed with a blue background on the Physical and Data Path Maps.

fabric port (F_port). A fabric port (F_port) that supports a node port (N_port) on a fibre-channel switch.

FAStT MSJ. See *FAStT Management Suite Java*.

FAStT Management Suite Java (FAStT MSJ) . A diagnostic and configuration utility that can be used on Linux, Microsoft Windows, and Novell NetWare host systems. In Linux, it is also used with the QLRemote agent to define preferred and non-preferred paths for logical drives.

feature enable identifier. A unique identifier for the storage subsystem, which is used in the process of generating a premium feature key. See also *premium feature key*.

fibre channel (FC). A bidirectional, full-duplex, point-to-point, serial data channel structured for high performance capability. Physically, fibre channel interconnects devices, such as host systems and servers, FC hubs and disk arrays, through ports, called node ports (N_ports), in one of three topologies: a point-to-point link, an arbitrated loop, or a cross point switched network, which is called a fabric. FC can interconnect two devices in a point-to-point topology, from two to 126 devices in an arbitrated loop. FC is a generalized transport mechanism that can transport any existing protocol, such as small computer system interface (SCSI), in FC frames.

Fibre Channel Protocol (FCP) for small computer system interface (SCSI). A high-level fibre-channel mapping layer (FC-4) that uses lower-level fibre-channel

(FC-PH) services to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target across the FC link by using FC frame and sequence formats.

field replaceable unit (FRU).   An assembly that is replaced in its entirety when any one of its components fails. In some cases, a FRU might contain other field replaceable units.

FRU.   See *field replaceable unit*.

GBIC.   See *gigabit interface converter*.

general purpose interface bus (GPIB).   An 8-bit parallel bus developed for the exchange of information between computers and industrial automation equipment.

gigabit interface converter (GBIC).   A transceiver that performs serial, optical-to-electrical and electrical-to-optical signal conversions for high-speed networking. A GBIC can be hot swapped. See also *small form-factor pluggable*.

GPIB.   See *general purpose interface bus*.

graphical user interface (GUI).   A type of computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution graphics, pointing devices, menu bars and other menus, overlapping windows, icons, and the object-action relationship.

GUI.   See *graphical user interface.*

HBA.   See *host bus adapter.*

hdisk.   An AIX term that represents a logical unit number (LUN) on an array.

host.   A system that is directly attached to the storage subsystem through a fibre-channel input/output path. This system is used to serve data (typically in the form of files) from the storage subsystem. A system can be both a storage management station and a host simultaneously.

host bus adapter (HBA).   An interface between the fibre-channel network and a workstation or server.

host computer.   See *host*.

host group.   A host group is 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.

host port.   Host ports physically reside on the host adapters and are automatically discovered by the FAStT Storage Manager software. To give a host computer access to a partition, you must define its associated host ports.

hot swap.   To replace a hardware component without turning off the system.

hub.   In a network, a point at which circuits are either connected or switched. For example, in a star network, the hub is the central node; in a star/ring network, it is the location of wiring concentrators.

IBMSAN driver.   The IBM Storage Area Network device driver that is used in a Novell NetWare environment to provide multipath I/O support to the storage controller.

IC.   See *integrated circuit.*

IDE.   See *integrated drive electronics.*

in-band.   Transmission of management protocol over the fibre-channel transport.

Industry Standard Architecture (ISA).   A bus standard for IBM compatibles that allows components to be added as cards plugged into standard expansion slots. ISA was originally introduced in the IBM PC/XT with an 8-bit data path. It was later expanded to permit a 16-bit data path when IBM introduced the PC/AT.

initial program load (IPL).   The part of the boot sequence during which a computer system copies the operating system kernel into main memory and runs it.

integrated circuit (IC).   A microelectronic semiconductor device that consists of many interconnected transistors and other components. ICs are constructed on a small rectangle cut from a silicon crystal or other semiconductor material. The small size of these circuits allows high speed, low power dissipation, and reduced manufacturing cost compared with board-level integration. Also known as a *chip*.

integrated drive electronics (IDE).   Also known as an Advanced Technology Attachment Interface (ATA). A disk drive interface based on the 16-bit IBM personal computer Industry Standard Architecture (ISA) in which the controller electronics reside on the drive itself, eliminating the need for a separate adapter card.

integrated services digital network (ISDN).   A digital end-to-end telecommunication network that supports multiple services including, but not limited to, voice and data. ISDNs are used in public and private network architectures.

interrupt request (IRQ).   A type of input found on many processors that causes the processor to suspend normal instruction execution temporarily and start executing an interrupt handler routine. Some processors have several interrupt request inputs that allow different priority interrupts.

Internet Protocol (IP) address. The unique 32-bit address that specifies the location of each device or workstation on the Internet. For example, 9.67.97.103 is an IP address.

IP address. See *Internet Protocol address.*

IPL. See *initial program load.*

IRQ. See *interrupt request.*

ISA. See *Industry Standard Architecture.*

ISDN. See *Integrated Services Digital Network.*

isolated group. A collection of isolated devices not connected to the storage area network (SAN) but discovered by the SANavigator tool. The isolated group displays with a gray background near the bottom of the Physical and Data Path Maps.

Java Runtime Environment (JRE). A subset of the Java Development Kit (JDK) for end users and developers who want to redistribute the JRE. The JRE consists of the Java virtual machine, the Java Core Classes, and supporting files.

JRE. See *Java Runtime Environment.*

label. A discovered or user entered property value that is displayed underneath each device in the Physical and Data Path Maps.

LAN. See *local area network.*

LBA. See *logical block addressing.*

local area network (LAN). A computer network located on a user's premises within a limited geographic area.

logical block addressing (LBA). A hard disk sector addressing scheme in which the addressing conversion is performed by the hard disk firmware. LBA is used on all small computer system interface (SCSI) hard disks and on ATA-2 conforming intergrated drive electronics (IDE) hard disks.

logical unit number (LUN). An identifier used on a small computer systems interface (SCSI) bus to distinguish among up to eight devices (logical units) with the same SCSI ID.

loop address. The unique ID of a node in fibre-channel loop topology sometimes referred to as a loop ID.

loop group. A collection of storage area network (SAN) devices that are interconnected serially in a single loop circuit. Loop groups are discovered by the SANavigator tool and displayed with a gray background on the Physical and Data Path Maps.

loop port (FL_port). A node port (N-port) or fabric port (F-port) that supports arbitrated loop functions associated with an arbitrated loop topology.

LUN. See *logical unit number.*

MAC. See *medium access control.*

medium access control (MAC). In local area networks (LANs), the sublayer of the data link control layer that supports medium-dependent functions and uses the services of the physical layer to provide services to the logical link control sublayer. The MAC sublayer includes the method of determining when a device has access to the transmission medium.

man pages. In UNIX-based operating systems, online documentation for operating system commands, subroutines, system calls, file formats, special files, stand-alone utilities, and miscellaneous facilities. Invoked by the **man** command.

management information base (MIB). The information that is on an agent. It is an abstraction of configuration and status information.

MCA. See *micro channel architecture.*

MIB. See *management information base.*

micro channel architecture (MCA). IBM's proprietary bus that is used in high-end PS/2 personal computers. Micro channel is designed for multiprocessing and functions as either a 16-bit or 32-bit bus. It eliminates potential conflicts that arise when installing new peripheral devices.

MIDI. See *musical instrument digital interface.*

model. The model identification assigned to a device by its manufacturer.

musical instrument digital interface (MIDI). A protocol that allows a synthesizer to send signals to another synthesizer or to a computer, or a computer to a musical instrument, or a computer to another computer.

NDIS. See *network device interface specification.*

network device interface specification (NDIS). An application programming interface (API) definition that allows disk operating systems (DOS) or OS/2 to support one or more network adapters and protocol stacks. NDIS is a 16-bit, Ring O (for the OS/2 operating system) application programming interface (API) that defines a specific way for writing drivers for layers 1 and 2 of the OSI model. NDIS also handles the configuration and binding of these network drivers to multiple protocol stacks.

network management station (NMS). In the Simple Network Management Protocol (SNMP), a station that

executes management application programs that monitor and control network elements.

NMI. See *non-maskable interrupt*.

NMS. See *network management station*.

non-maskable interrupt (NMI). A hardware interrupt that another service request cannot overrule (mask). An NMI bypasses and takes priority over interrupt requests generated by software, the keyboard, and other such devices and is issued to the microprocessor only in disastrous circumstances, such as severe memory errors or impending power failures.

node. A physical device that allows for the transmission of data within a network.

node port (N_port). A fibre-channel defined hardware entity that performs data communications over the fibre-channel link. It is identifiable by a unique worldwide name. It can act as an originator or a responder.

N_port. See *node port*.

nonvolatile storage (NVS). A storage device whose contents are not lost when power is cut off.

NVS. See *nonvolatile storage*.

NVSRAM. Nonvolatile storage random access memory. See *nonvolatile storage*.

Object Data Manager (ODM). An AIX proprietary storage mechanism for ASCII stanza files that are edited as part of configuring a drive into the kernel.

ODM. See *Object Data Manager*.

out-of-band. Transmission of management protocols outside of the fibre-channel network, typically over Ethernet.

PCI local bus. See *peripheral component interconnect local bus*.

PDF. See *portable document format*.

peripheral component interconnect local bus (PCI local bus). A standard that Intel Corporation introduced for connecting peripherals. The PCI local bus allows up to 10 PCI-compliant expansion cards to be installed in a computer at a time. Technically, PCI is not a bus but a bridge or mezzanine. It runs at 20 - 33 MHz and carries 32 bits at a time over a 124-pin connector or 64 bits over a 188-pin connector. A PCI controller card must be installed in one of the PCI-compliant slots. The PCI local bus is processor independent and includes buffers to decouple the central processing unit (CPU) from relatively slow peripherals, allowing them to operate asynchronously. It also allows for multiplexing, a technique that permits more than one electrical signal to be present on the PCI local bus at a time.

performance events. Events related to thresholds set on storage area network (SAN) performance.

polling delay. The time in seconds between successive discovery processes during which discovery is inactive.

port. The hardware entity that connects a device to a fibre-channel topology. A device can contain one or more ports.

portable document format (PDF). A standard specified by Adobe Systems, Incorporated, for the electronic distribution of documents. PDF files are compact; can be distributed globally by e-mail, the Web, intranets, or CD-ROM; and can be viewed with the Acrobat Reader, which is software from Adobe Systems that can be downloaded at no cost from the Adobe Systems home page.

premium feature key. A file that the storage subsystem controller uses to enable an authorized premium feature. The file contains the feature enable identifier of the storage subsystem for which the premium feature is authorized, and data about the premium feature. See also *feature enable identifier*.

private loop. A freestanding arbitrated loop with no fabric attachment.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of the program.

PTF. See *program temporary fix*.

RAM. See *random-access memory*.

random-access memory (RAM). A temporary storage location in which the central processing unit (CPU) stores and executes its processes.

RDAC. See *redundant disk array controller*.

read-only memory (ROM). Memory in which the user cannot changed stored data except under special conditions.

recoverable virtual shared disk (RVSD). A virtual shared disk on a server node configured to provide continuous access to data and file systems in a cluster.

red, green, blue (RGB). (1) Color coding in which the brightness of the additive primary colors of light, red, green, and blue are specified as three distinct values of white light. (2) Pertaining to a color display that accepts signals that represent red, green, and blue.

redundant disk array controller (RDAC). (1) In hardware, a redundant set of controllers (either active/passive or active/active). (2) In software, a layer that manages the input/output (I/O) through the active controller during normal operation and transparently

reroutes I/Os to the other controller in the redundant set if a controller or I/O path fails.

RGB. See *red, green, blue*.

ROM. See *read-only memory*.

router. A computer that determines the path of network traffic flow. The path selection is made from several paths based on information obtained from specific protocols, algorithms that attempt to identify the shortest or best path, and other criteria such as metrics or protocol-specific destination addresses.

RVSD. See *recoverable virtual shared disk*.

SAN. See *storage area network*.

scope. Defines a group of controllers by their Internet Protocol (IP) addresses. You must create and configure a scope so that dynamic IP addresses can be assigned to controllers on your network.

SCSI. See *small computer system interface*.

segmented loop port (SL_port). A segmented loop port (SL_port) allows you to divide a fibre-channel private loop into multiple segments. Each segment can pass frames around as an independent loop and can connect through the fabric to other segments of the same loop.

sense data. (1) Data sent with a negative response, indicating the reason for the response. (2) Data describing an I/O error. Sense data is presented to a host system in response to a sense request command.

serial storage architecture (SSA). An interface specification from IBM in which devices are arranged in a ring topology. SSA, which is compatible with small computer system interface (SCSI) devices, allows full-duplex packet multiplexed serial data transfers at rates of 20 Mbps in each direction.

server. A functional hardware and software unit that delivers shared resources to workstation client units on a computer network.

server/device events. Events that occur on the server or a designated device that meet criteria that the user sets.

SFP. See *small form-factor pluggable*.

Simple Network Management Protocol (SNMP). In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

SL_port. See *segmented loop port*.

SMagent. The FAStT Storage Manager optional Java-based host-agent software, which can be used on Microsoft Windows, Novell Netware, HP-UX and Solaris host systems to manage storage subsystems through the host fibre-channel connection.

SMclient. The FAStT Storage Manager client software, which is a Java-based graphical user interface (GUI) that is used to configure, manage, and troubleshoot storage servers and expansion units in a FAStT storage subsystem. SMclient can be used on a host system or on a storage management station.

SMruntime. A Java compiler for the SMclient.

SMutil. The FAStT Storage Manager utility software that is used on Microsoft Windows, HP-UX and Solaris host systems to register and map new logical drives to the operating system. In Microsoft Windows, it also contains a utility to flush the system's OS cached data for a particular drive before creating a FlashCopy.

small computer system interface (SCSI). A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

small form-factor pluggable (SFP). An optical transceiver that is used to convert signals between optical fiber cables and switches. An SFP is smaller than a gigabit interface converter (GBIC). See also *gigabit interface converter*.

SNMP. See *Simple Network Management Protocol*.

SNMPv1. The original standard for Simple Network Management Protocol (SNMP) is now referred to as SNMPv1, as opposed to SNMPv2, a revision of SNMP. See also *Simple Network Management Protocol*.

SNMP time-out. The maximum amount of time the SANavigator tool will wait for a device to respond to a request. The specified time applies to one retry only.

SNMP trap events. Simple Network Management Protocol (SNMP) is based on a manager/agent model. SNMP includes a limited set of management commands and responses. The management system issues messages that tell an agent to retrieve various object variables. The managed agent sends a Response message to the management system. That message is an event notification, called a trap, that identifies conditions, such as thresholds, that exceed a predetermined value.

SRAM. See *static random access memory*.

SSA. See *serial storage architecture*.

static random access memory (SRAM). Random access memory based on the logic circuit known as flip-flop. It is called *static* because it retains a value as long as power is supplied, unlike dynamic random access memory (DRAM), which must be regularly

refreshed. It is however, still volatile, meaning that it can lose its contents when the power is switched off.

storage area network (SAN). A network that links servers or workstations to disk arrays, tape backup subsystems, and other devices, typically over fibre channel.

storage management station. A system that is used to manage the storage subsystem. A storage management station does not need to be attached to the storage subsystem through the fibre-channel input/output path.

storage partition. Storage partitions are 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. In the FAStT Storage Manager client, the Topology view of the Mappings window displays the default host group, the defined host group, the 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.

subnet. An interconnected but independent segment of a network that is identified by its Internet Protocol (IP) address.

super video graphics array (SVGA). A video display standard that Video Electronics Standards Association (VESA) created to provide high resolution color display on IBM PC compatible personal computers. The resolution is 800 x 600 4-bit pixels. Each pixel can therefore be one of 16 colors.

SVGA. See *super video graphics array*.

sweep method. A method of sending Simple Network Management Protocol (SNMP) requests for information to all the devices on a subnet by sending the request to every device on the network. Sweeping an entire network can take a half an hour or more. If broadcast is disabled, the recommended method is to enter the individual Internet Protocol (IP) addresses of the storage area network (SAN) devices into the SANavigator tool. This method produces good results without unnecessarily using time to wait for responses from every IP address in the subnet, especially for IP addresses where no devices are present. There might, however, be times when a full subnet sweep will produce valuable diagnostic information about the network or a device's configuration.

switch. A fibre-channel device that provides full bandwidth per port and high-speed routing of data by using link-level addressing.

switch group. A switch and the collection of devices connected to it that are not in other groups. Switch

groups are discovered by the SANavigator tool and displayed with a gray background on the Physical and Data Path Maps.

system name. Device name assigned by the vendor's third-party software.

TCP. See *Transmission Control Protocol*.

TCP/IP. See *Transmission Control Protocol/Internet Protocol*.

terminate and stay resident program (TSR program). A program that installs part of itself as an extension of DOS when it is executed.

TFT. See *thin-film transistor*.

thin-film transistor (TFT). A transistor created by using thin film methodology.

topology. The physical or logical arrangement of devices on a network. The three fibre-channel topologies are fabric, arbitrated loop, and point-to-point. The default topology for the disk array is arbitrated loop.

TL_port. See *translated loop port*.

transceiver. A device that is used is used to transmit and receive data. "Transceiver" is an abbreviation of "transmitter-receiver."

translated loop port (TL_port). Each translated loop port (TL_port) connects to a private loop and allows connectivity between the private loop devices and *off loop* devices (devices not connected to that particular TL_port).

Transmission Control Protocol (TCP). A communication protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packed-switched communication networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP). A set of communication protocols that provide peer-to-peer connectivity functions for both local and wide-area networks.

trap. In the Simple Network Management Protocol (SNMP), a message sent by a managed node (agent function) to a management station to report an exception condition.

trap recipient. Receiver of a forwarded Simple Network Management Protocol (SNMP) trap. Specifically, a trap receiver is defined by an Internet Protocol (IP) address and port to which traps are sent. Presumably, the actual recipient is a software application running at the IP address and listening to the port.

TSR program. See *terminate and stay resident program*.

uninterruptible power supply. A source of power from a battery that is installed between a computer system and its power source. The uninterruptible power supply keeps the system running if a commercial power failure occurs, until an orderly shutdown of the system can be performed.

user action events. Actions that the user takes, such as changes in the storage area network (SAN), changed settings, and so on. Each such action is considered a *user action event*.

vendor. Property value that the SANavigator tool uses to launch third-party software. Vendor property might be discovered but will always remain editable.

VGA. See *video graphics adapter*.

video graphics adapter (VGA). A computer adapter that provides high-resolution graphics and a total of 256 colors.

video random access memory (VRAM). A special type of dynamic random access memory (DRAM) used in high-speed video applications, designed for storing the image to be displayed on a computer's monitor.

VRAM. See *video random access memory*.

WORM. See *write-once read-many*.

worldwide name (WWN). A registered, unique 64–bit identifier assigned to nodes and ports.

write-once read-many (WORM). Any type of storage medium to which data can be written only a single time, but can be read from any number of times. After the data is recorded, it cannot be altered. Typically the storage medium is an optical disk whose surface is permanently etched by using a laser in order to record information. WORM media are high-capacity storage devices and have a significantly longer shelf life than magnetic media.

WWN. See *worldwide name*.

XGA. See *eXtended graphics array*.

zoning. A function that allows segmentation of nodes by address, name, or physical port and is provided by fabric switches or hubs.

Index

Readers' comments — we would like to hear from you.

IBM TotalStorage FAStT Storage Manager
Concepts Guide

Publication No. GC26-7560-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you?   ☐ Yes   ☐ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.

IBM®

Fold and Tape   **Please do not staple**   Fold and Tape
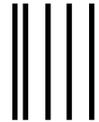
BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 40   ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
Dept. M86/Bldg.050-3
5600 Cottle Road
San Jose, CA
U.S.A. 95193-0001

Fold and Tape   **Please do not staple**   Fold and Tape

IBM®
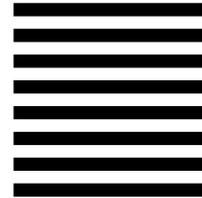
Part Number: 25P1661

Printed in USA

(1P) P/N: 25P1661