

IBM TotalStorage DS4000 Storage Manager Version 9



Concepts Guide

IBM TotalStorage DS4000 Storage Manager Version 9



Concepts Guide

Note:

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 129.

Third Edition (September 2005)

© Copyright International Business Machines Corporation 2005. All rights reserved.

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

Contents

Figures	vii
Tables	ix
About this document	xi
Who should read this document	xi
DS4000 installation task overview	xi
How this document is organized	xiv
Notices that this document uses	xv
Getting information, help, and service	xv
Before you call	xv
Using the documentation	xv
Web sites	xvi
Software service and support	xvii
Hardware service and support	xvii
How to send your comments	xvii
Chapter 1. Introduction	1
Storage Manager documentation and readme files	1
Product updates	1
FAST product renaming	2
Machine types and supported software	3
Terms to know	5
New features and enhancements	6
FAST product renaming	7
Controller firmware 6.14 and 6.15: New features	7
Controller firmware 6.12: New features	7
Controller firmware 6.10: New features	8
Storage Manager premium features	10
Storage subsystem components	10
Storage subsystem model types	11
Storage partitioning specifications	12
Software components	14
Host types	18
System requirements	19
Hardware requirements	19
Storage subsystem management	21
Direct (out-of-band) management method	21
Host-agent (in-band) management method	23
Reviewing a sample network	25
Managing coexisting storage subsystems	26
Managing the storage subsystem using the graphical user interface	27
Enterprise Management window	27
Populating a management domain	28
Subsystem Management window	29
Subsystem Management window tabs	30
The Subsystem Management window menus	31
The script editor	35
Using the script editor	36
Adding comments to a script	37
The command line interface (SMcli)	38
Using SMcli	38
Command line interface parameters	39

Usage and formatting requirements	42
SMcli examples.	43
Chapter 2. Storing and protecting your data	45
Logical drives	45
Dynamic Logical Drive Expansion	47
Arrays	47
Dynamic Capacity Expansion	48
Fibre-channel I/O data path failover support	48
Auto-Logical Drive Transfer feature	48
Redundant disk array controller (RDAC)	49
Operating system specific failover protection	49
Default settings for failover protection	50
Redundant array of independent disks (RAID)	51
Protecting data in the controller cache memory	53
Configuring hot-spare drives	54
Media scan	55
Errors reported by a media scan	57
Media scan settings	58
Media scan duration	58
Copy services and the DS4000 Storage Subsystem	59
FlashCopy	61
VolumeCopy	62
Copying data for greater access	62
Backing up data	62
Restoring FlashCopy logical drive data to the base logical drive	62
Enhanced Remote Mirroring option	63
Enhanced Remote Mirroring option enhancements	63
Logical drives on a remote mirror setup	65
Write modes	66
Mirror relationships	66
Managing Persistent Reservations	67
Configuring storage subsystem password protection	67
Chapter 3. Configuring storage subsystems	69
Creating logical drives	69
Storage partitioning	70
Switch zoning	70
Storage partitioning terminology.	70
Obtaining a feature key.	72
Heterogeneous hosts overview	72
Chapter 4. Maintaining and monitoring storage subsystems	75
Using the Task Assistant	75
Maintaining storage subsystems in a management domain.	78
Storage subsystem status quick reference	78
Failure notification.	79
Updating the firmware in the storage subsystem and storage expansion enclosures.	79
Downloading controller firmware	80
Traditional controller firmware download.	80
The staged controller firmware download feature	81
Downloading NVSRAM	81
Downloading NVSRAM from a firmware image	81
Downloading NVSRAM as a standalone image	81
Downloading drive firmware	81

General Considerations	82
Parallel drive firmware download	82
Environmental services module card	83
Downloading ESM firmware	83
Viewing and recovering missing logical drives	83
Alert notification overview	84
Configuring mail server and sender address	85
Selecting the node for notification	85
Setting alert destinations	85
Configuring alert destinations for storage subsystem critical-event notification	85
Event Monitor overview	86
Installing the Event Monitor	86
Setting alert notifications	86
Synchronizing the Enterprise Management window and Event Monitor	87
Recovery Guru	88
Chapter 5. Tuning storage subsystems	93
Balancing the Fibre Channel I/O load	93
Optimizing the I/O transfer rate	94
Optimizing the Fibre Channel I/O request rate	94
Determining the Fibre Channel I/O access pattern and I/O size	94
Enabling write-caching	95
Optimizing the cache-hit percentage	95
Choosing appropriate RAID levels	95
Choosing an optimal logical-drive modification priority setting	95
Choosing an optimal segment size	96
Defragmenting files to minimize disk access	96
Chapter 6. Critical event problem solving	97
Appendix A. Online help task reference	113
Populating a management domain	113
Configuring storage subsystems	114
Using the Script Editor	114
Configuring storage partitions	115
Protecting data	116
Event notification	116
Recovering from problems	117
Miscellaneous system administration	117
Security	118
Performance and tuning	118
Appendix B. Additional DS4000 documentation	119
DS4500 Fibre Channel Storage Server library	119
DS4400 Fibre Channel Storage Server library	120
DS4300 Fibre Channel Storage Server library	121
DS4100 SATA Storage Server library	122
DS4800 Storage Subsystem library	123
Other DS4000 and DS4000-related documents	124
Appendix C. Accessibility	127
Notices	129
Trademarks	129
Important notes	130

Glossary	131
Index	141

Figures

1. Direct (out-of-band) managed storage subsystems	23
2. Host-agent (in-band) managed storage subsystems	25
3. Sample network using direct and host-agent managed storage subsystems	26
4. The Enterprise Management window	27
5. Device tree with a management domain	28
6. Subsystem Management window Logical View and Physical View	30
7. The script editor window	36
8. Unconfigured and free capacity nodes	69
9. The task assistant in the Enterprise Management window	76
10. The task assistant in the Subsystem Management window	77
11. Monitoring storage subsystem health using the Enterprise Management window	78
12. Event monitoring example	87
13. Location of the Recovery Guru toolbar button	89
14. Recovery Guru window	89
15. Recovery Guru window showing Replaced status icon	90
16. Recovery Guru window showing Optimal status	91

Tables

1.	Where to find DS4000 installation and configuration procedures	xii
2.	Mapping of FAST names to DS4000 Series names	3
I 3.	Machine types, supported controller firmware versions, and supported Storage Manager software	4
4.	Old and new terminology	5
5.	Storage subsystem physical components	11
6.	Storage subsystem logical components	12
I 7.	Storage partitioning specifications per DS4000 storage subsystem model	13
8.	Storage Manager software components	14
9.	Storage management architecture hardware components	20
10.	Default settings for controllers with firmware version 05.00.xx or later	22
11.	Subsystem Management window tabs	31
12.	The Subsystem Management window menus	32
13.	Command line parameters	39
14.	RAID level configurations	52
15.	Errors discovered during a media scan	57
16.	Restrictions to copy services premium feature support	60
17.	Storage partitioning terminology	70
18.	Storage subsystem status icon quick reference	78
19.	Performance Monitor tuning options in the Subsystem Management window	93
20.	Critical events	97
21.	TotalStorage DS4500 Fibre Channel Storage Server document titles by user tasks	119
22.	TotalStorage DS4400 Fibre Channel Storage Server document titles by user tasks	120
23.	TotalStorage DS4300 Fibre Channel Storage Server document titles by user tasks	121
24.	TotalStorage DS4100 SATA Storage Server document titles by user tasks	122
25.	TotalStorage DS4800 Storage Subsystem document titles by user tasks	123
26.	TotalStorage DS4000 and DS4000-related document titles by user tasks	124
27.	DS4000 Storage Manager alternate keyboard operations	127

About this document

This document provides the conceptual framework to help you understand the IBM® TotalStorage® DS4000 Storage Manager Version 9.1x for the following operating-system environments:

- Windows® 2000
- Windows® Server 2003
- Novell NetWare
- Sun Solaris
- Hewlett-Packard HP-UX
- IBM AIX®
- Red Hat and SUSE Linux
- Red Hat and SUSE Linux on POWER
- 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.

DS4000 installation task overview

Table 1 on page xii provides a sequential list of many installation and configuration tasks that are common to most DS4000 configurations. When you install and configure your DS4000 storage subsystem, refer to this table to find the documentation that explains how to complete each task.

See also: The *DS4000 Storage Server and Storage Expansion Enclosure Quick Start Guide* provides an excellent overview of the installation process.

Table 1. Where to find DS4000 installation and configuration procedures

	Installation task	Where to find information or procedures
1	Plan the installation	<ul style="list-style-type: none"> • DS4000 Storage Manager Concepts Guide • DS4000 Storage Manager Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER • DS4000 Storage Manager Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux • DS4100 Storage Server Installation, User's, and Maintenance Guide • DS4300 Fibre Channel Storage Server Installation and User's Guide • DS4400 Fibre Channel Storage Server Installation and Support Guide • DS4500 Fibre Channel Storage Server Installation and Support Guide • DS4800 Storage Subsystem Installation, User's, and Maintenance Guide
2	Mount the DS4000 storage subsystem in the rack	<ul style="list-style-type: none"> • DS4800 Installation, User's, and Maintenance Guide • DS4400 and DS4500 Rack Mounting Instructions • DS4300 Rack Mounting Instructions • DS4100 Installation, User's and Maintenance Guide
3	Mount the DS4000 EXP storage expansion unit in the rack	<ul style="list-style-type: none"> • DS4000 EXP100 Storage Expansion Unit Installation, User's and Maintenance Guide • DS4000 EXP700 and EXP710 Storage Expansion Enclosures Installation, User's, and Maintenance Guide • FASTT EXP500 Installation and User's Guide
4	Route the storage expansion unit Fibre Channel cables	<ul style="list-style-type: none"> • DS4100 Fibre Channel Cabling Instructions • DS4300 Fibre Channel Cabling Instructions • DS4400 Fibre Channel Cabling Instructions • DS4500 Fibre Channel Cabling Instructions • DS4800 Installation, User's, and Maintenance Guide
5	Route the host server Fibre Channel cables	<ul style="list-style-type: none"> • DS4100 Fibre Channel Cabling Instructions • DS4300 Fibre Channel Cabling Instructions • DS4400 Fibre Channel Cabling Instructions • DS4500 Fibre Channel Cabling Instructions • DS4800 Installation, User's, and Maintenance Guide
6	Power up the subsystem	<ul style="list-style-type: none"> • DS4100 Storage Server Installation, User's, and Maintenance Guide • DS4300 Fibre Channel Storage Server Installation and User's Guide • DS4400 Fibre Channel Storage Server Installation and Support Guide • DS4500 Fibre Channel Storage Server Installation and Support Guide • DS4800 Storage Subsystem Installation, User's, and Maintenance Guide

Table 1. Where to find DS4000 installation and configuration procedures (continued)

	Installation task	Where to find information or procedures
7	Configure DS4000 network settings	<ul style="list-style-type: none"> • DS4100 Storage Server Installation, User's, and Maintenance Guide • DS4300 Fibre Channel Storage Server Installation and User's Guide • DS4400 Fibre Channel Storage Server Installation and Support Guide • DS4500 Fibre Channel Storage Server Installation and Support Guide • DS4800 Storage Subsystem Installation, User's, and Maintenance Guide
8	Zone the fabric switch (SAN-attached only)	<ul style="list-style-type: none"> • DS4000 Storage Manager Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER • DS4000 Storage Manager Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux • DS4000 Storage Manager Copy Services Guide (describes switch zoning for the Remote Mirror Option) • See also the documentation provided by the switch manufacturer
9	Install DS4000 Storage Manager software on the management station	<ul style="list-style-type: none"> • DS4000 Storage Manager Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER • DS4000 Storage Manager Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux • DS4000 Storage Manager online help (for post-installation tasks)
10	Install host software (failover drivers) on host server	
11	Start DS4000 Storage Manager	
12	Set the DS4000 Storage Manager clock	
13	Set the DS4000 Storage Manager host default type	
14	Verify DS4000 subsystem health	<ul style="list-style-type: none"> • DS4100 Storage Server Installation, User's, and Maintenance Guide • DS4300 Fibre Channel Storage Server Installation and User's Guide • DS4400 Fibre Channel Storage Server Installation and Support Guide • DS4500 Fibre Channel Storage Server Installation and Support Guide • DS4800 Storage Subsystem Installation, User's, and Maintenance Guide

Table 1. Where to find DS4000 installation and configuration procedures (continued)

	Installation task	Where to find information or procedures
15	Enable DS4000 Storage Manager premium feature keys	<p>Copy Services premium features DS4000 Storage Manager Copy Services Guide</p> <p>FC/SATA Intermix premium feature DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview</p> <p>Storage Partitioning (and general premium features information)</p> <ul style="list-style-type: none"> • DS4000 Storage Manager Concepts Guide • DS4000 Storage Manager Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER • DS4000 Storage Manager Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux
16	Configure arrays and logical drives	<ul style="list-style-type: none"> • DS4000 Storage Manager Installation and Support Guide for AIX, HP-UX, Solaris and Linux on POWER
17	Configure host partitions	<ul style="list-style-type: none"> • DS4000 Storage Manager Installation and Support Guide for Windows 2000/Server 2003, NetWare, ESX Server, and Linux
18	Verify host access to DS4000 storage	<ul style="list-style-type: none"> • DS4000 Storage Manager online help

How this document is organized

Chapter 1, “Introduction,” on page 1 provides an introduction of the IBM DS4000 Storage Manager Version 9.1x, compares storage subsystem management methods, and describes the DS4000 Storage Manager Enterprise Management and Subsystem Management windows.

Chapter 2, “Storing and protecting your data,” on page 45 describes the various data protection features of the DS4000 Storage Subsystem. These features include input/output (I/O) data path failover support, Media Scan, and copy services.

Chapter 3, “Configuring storage subsystems,” on page 69 describes the frequently-used functions of the IBM DS4000 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 75, 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 93 discusses the tuning options that are available in IBM DS4000 Storage Manager Version 9.1x.

Chapter 6, “Critical event problem solving,” on page 97 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 113 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 DS4000 documentation,” on page 119 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, ensure that you take these steps to try to solve the problem yourself:

- Check all cables to ensure that they are connected.
- Check the power switches to ensure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system.
- To check for technical information, hints, tips, and new device drivers, go to the IBM Support Web site at:
www.ibm.com/pc/support/
- 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.

Web sites

DS4000 Storage Subsystems

Find the latest information about IBM TotalStorage disk storage systems, including all of the DS4000 Storage Subsystems:

www-1.ibm.com/servers/storage/disk/ds4000/

IBM TotalStorage products

Find information about all IBM TotalStorage products:

www.storage.ibm.com/

Support for IBM TotalStorage disk storage systems

Find links to support pages for all IBM TotalStorage disk storage systems, including DS4000 storage subsystems and expansion units:

www-1.ibm.com/servers/storage/support/disk/

TotalStorage DS4000 interoperability matrix

Find the latest information about operating system and HBA support, clustering support, storage area network (SAN) fabric support, and DS4000 Storage Manager feature support:

www-1.ibm.com/servers/storage/disk/ds4000/interop-matrix.html

DS4000 Storage Manager readme files

Find the latest Storage Manager readme files for each host operating system at the IBM TotalStorage support page by completing these steps:

1. Go to the following Web site:

www-1.ibm.com/servers/storage/support/disk/

2. Click the link for your storage subsystem.
3. When the subsystem page opens, click the **Download** tab.
4. When the download page opens, click the **Storage Mgr** tab then click on the appropriate link under the Current Versions and Readmes column.

Storage Area Network (SAN) support

Find information about using SAN switches, including links to user guides and other documents:

www.ibm.com/servers/storage/support/san/index.html

DS4000 technical support

Find downloads, hints and tips, documentation, parts information, HBA and Fibre Channel support:

www-1.ibm.com/servers/storage/support/disk/

Premium feature activation

Enable a premium feature on a DS4000 storage subsystem by using the online tool:

www-912.ibm.com/PremiumFeatures/jsp/keyInput.jsp

IBM publications center

Find IBM publications:

www.ibm.com/shop/publications/order/

IBM xSeries product support

Find up-to-date information about IBM xSeries products, services, and support:

www.ibm.com/eserver/xseries/

IBM personal computing support

Find service information for your IBM products, including supported options:

www.ibm.com/pc/support/

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/

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

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

International Business Machines Corporation
Department GZW
9000 South Rita Road
Tucson, AZ
85744-0001
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® DS4000 Storage Manager Version 9.1x.

Storage Manager documentation and readme files

Before you install the IBM DS4000 Storage Manager software, consult the following documentation:

Storage Manager readme files

Read these first.

1. For the most recent Storage Manager readme files for your operating system, see the following Web site:

www-1.ibm.com/servers/storage/support/disk/

2. Click the link for your storage subsystem.
3. When the subsystem page opens, click the **Download** tab.
4. When the download page opens, click the **Storage Mgr** tab then click on the appropriate link under the Current Versions and Readmes column.

Important: Updated readme files contain the latest device driver versions, firmware levels and other information that supersedes this document.

IBM DS4000 Storage Manager Installation and Support Guides

Use the installation and support guide for your operating system or platform to set up, install, configure, and work with the IBM DS4000 Storage Manager Version 9.1x.

After you complete all of the Storage Manager and host installation procedures, refer to the following online help systems, which contain information and procedures that are common to all host 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.

Note: To access the help systems from the Enterprise Management and Subsystem Management windows in IBM DS4000 Storage Manager Version 9.1x, click **Help** on the toolbar, or press F1.

Product updates

Important

In order to keep your system up to date with the latest firmware and other product updates, use the information below to register and use the **My support** web site.

Download the latest versions of the DS4000 Storage Manager host software, DS4000 storage server controller firmware, DS4000 drive expansion enclosure ESM firmware, and drive firmware at the time of the initial installation and when product updates become available.

To be notified of important product updates, you must first register at the IBM Support and Download Web site:

www-1.ibm.com/servers/storage/support/disk/index.html

In the **Additional Support** section of the Web page, click **My support**. On the next page, if you have not already done so, register to use the site by clicking **register now**.

Perform the following steps to receive product updates:

1. After you have registered, type your user ID and password to log into the site. The **My support** page opens.
2. Click **add products**. A pull-down menu displays.
3. In the pull-down menu, select **Storage**. Another pull-down menu displays.
4. In the new pull-down menu, and in the subsequent pull-down menus that display, select the following topics:
 - **Computer Storage**
 - **Disk Storage Systems**
 - **TotalStorage DS4000 Midrange Disk Systems & FAStT Stor Svrs**

Note: During this process a check list displays. Do not check any of the items in the check list until you complete the selections in the pull-down menus.

5. When you finish selecting the menu topics, place a check in the box for the machine type of your DS4000 series product, as well as any other attached DS4000 series product(s) for which you would like to receive information, then click **Add products**. The **My support** page opens again.
6. On the **My support** page, click the **Edit profile** tab, then click **Subscribe to email**. A pull-down menu displays.
7. In the pull-down menu, select **Storage**. A check list displays.
8. Place a check in each of the following boxes:
 - a. **Please send these documents by weekly email**
 - b. **Downloads and drivers**
 - c. **Flashes**and any other topics that you may be interested in, and then click **Update**.
9. Click **Sign out** to log out of My Support.

FAStT product renaming

IBM has renamed some FAStT family products. Table 2 on page 3 identifies each DS4000 product name with its corresponding previous FAStT product name. Note that this change of **product name only** indicates no change in functionality or warranty. All products listed below with new names are functionally-equivalent and fully-interoperable. Each DS4000 product retains full IBM service as outlined in service contracts issued for analogous FAStT products.

Table 2. Mapping of FAStT names to DS4000 Series names

Previous FAStT Product Name	Current DS4000 Product Name
IBM TotalStorage FAStT Storage Server	IBM TotalStorage DS4000
FAStT	DS4000
FAStT Family	DS4000 Mid-range Disk System
FAStT Storage Manager vX.Y (for example v9.10)	DS4000 Storage Manager vX.Y (for example v9.10)
FAStT100	DS4100
FAStT600	DS4300
FAStT600 with Turbo Feature	DS4300 Turbo
FAStT700	DS4400
FAStT900	DS4500
EXP700	DS4000 EXP700
EXP710	DS4000 EXP710
EXP100	DS4000 EXP100
FAStT FlashCopy	FlashCopy for DS4000
FAStT VolumeCopy	VolumeCopy for DS4000
FAStT Remote Mirror (RM)	Enhanced Remote Mirroring for DS4000
FAStT Synchronous Mirroring	Metro Mirroring for DS4000
	Global Copy for DS4000 (New Feature = Asynchronous Mirroring without Consistency Group)
	Global Mirroring for DS4000 (New Feature = Asynchronous Mirroring with Consistency Group)

Machine types and supported software

Table 3 on page 4 provides a list of machine types and supported storage management software.

Notes:

1. Controller firmware versions 06.14.xx.xx and 06.15.xx.xx support DS4800 storage subsystems only. All other DS4000 Storage Subsystems continue to run firmware versions prior to 6.14.xx.xx.
2. Storage subsystem controller firmware must be at version 4.01.xx.xx, or later (05.xx.xx.xx or later with Windows host servers) to be managed by Storage Manager Version 9.1x. The only Storage Manager version 9.1x function that you can perform on a storage subsystem with controller firmware earlier than 04.01.xx.xx is download firmware in order to upgrade the controller firmware to a firmware version later than 4.01.xx.xx.

To ensure the highest level of compatibility and error-free operation, ensure that the controller firmware for your DS4000 Storage Subsystem is the latest firmware version for the storage subsystem model. In Table 3 on page 4, the latest client code software and controller firmware versions are indicated in **bold** text.

Table 3. Machine types, supported controller firmware versions, and supported Storage Manager software

Product name	Machine type	Model	Controller firmware version	Supported storage manager software version
IBM TotalStorage® DS4800 Storage Subsystem	1815	82A/H 84A/H 88A/H	06.14.xx.xx, 06.15.xx.xx	9.14, 9.15
IBM TotalStorage DS4100 Storage Subsystem (Base Model)	1724	100	6.10.xx.xx, 06.12.xx.xx	8.42, 9.10, 9.12, 9.14, 9.15
IBM TotalStorage DS4100 Storage Subsystem (Single Controller Model)	1724	1SC 1S	5.42.xx.xx	
IBM TotalStorage DS4500 Disk Storage Subsystem	1742	90X 90U	5.30.xx.xx, 5.40.xx.xx, 5.41.xx.xx (supports EXP100 only), 6.10.xx.xx, 06.12.xx.xx	8.3, 8.4, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15
IBM TotalStorage DS4400 Disk Storage Subsystem	1742	1RU 1RX	5.00.xx.xx, 5.20.xx.xx, 5.21.xx.xx, 5.30.xx.xx, 5.40.xx.xx, 6.10.xx.xx, 6.12.xx.xx	8.0, 8.2, 8.21, 8.3, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15
IBM TotalStorage DS4300 Disk Storage Subsystem (Single Controller)	1722	6LU 6LX	5.34.xx.xx	8.41.xx.03 or later, 8.42, 9.10, 9.12, 9.14, 9.15
IBM TotalStorage DS4300 Disk Storage Subsystem (Base Model)		60U 60X	5.33.xx.xx, 5.34.xx.xx, 5.40.xx.xx, 6.10.xx.xx, 6.12.xx.xx	8.3, 8.4, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15
IBM TotalStorage DS4300 Disk Storage Subsystem (Turbo Model)		60U 60X	5.41.xx.xx (supports EXP100 only), 6.10.xx.xx, 6.12.xx.xx	
IBM Netfinity® FAST500 RAID Controller Enclosure Unit (no longer available for purchase)	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, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15
IBM FAST200 High Availability (HA) Storage Subsystem (no longer available for purchase)	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, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15
IBM FAST200 Storage Subsystem (no longer available for purchase)	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, 8.41, 8.42, 9.10, 9.12, 9.14, 9.15

Table 3. Machine types, supported controller firmware versions, and supported Storage Manager software (continued)

Product name	Machine type	Model	Controller firmware version	Supported storage manager software version
IBM Netfinity Fibre Channel RAID Controller Unit (no longer available for purchase)	3526	1RU 1RX	4.x	7.0, 7.01, 7.02, 7.10, 8.0, 8.2, 8.21, 8.3, 8.41, 8.42, 9.10, 9.14, 9.15

Notes:

- All of the controller firmware versions listed in the table are available free-of-charge.
- Storage subsystems with controller firmware version 04.00.02.xx through 4.10.xx.xx must be managed with Storage Manager 8.x.
- Controller firmware level 06.12.xx.xx supports EXP100 SATA expansion enclosures with the following storage subsystems:
 - DS4100 and DS4300 Base models
 - DS4300 Turbo models
 - DS4400
 - DS4500

If you want to upgrade to 06.12.xx.xx and your controller firmware level is currently 05.41.1x.xx, you must first upgrade to firmware version 05.41.5x.xx (provided on the CD that is shipped with the EXP100.) After your firmware is at level 05.41.5x.xx, you can then upgrade to 06.12.xx.xx.

- Firmware levels 5.40.xx.xx and earlier provide support for EXP500 and EXP700 storage expansion enclosures only. For EXP710 support, firmware versions 06.1x.xx.xx or later are required.

Note: Firmware versions 06.10.11.xx and later support intermixing Fibre Channel and SATA storage expansion enclosures in the same DS4000 storage subsystem, if the DS4000 FC/SATA Intermix premium feature option is purchased. Contact your IBM representatives or reseller for more information.

See “Machine types and supported software” on page 3 for the latest firmware versions that are available for specific DS4000 Storage Subsystem models.

Terms to know

If you are upgrading from a previous version of Storage Manager, 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 4 provides a list of some of the old and new terms.

Table 4. Old and new terminology

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

Table 4. Old and new terminology (continued)

Term used in previous versions	New term
Controller module	Controller enclosure
Environmental card CRU	Environmental service module (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
Volume group	Array

Note: In Storage Manager 7.10 and later, 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

DS4000 Storage Manager 9.15 supports controller firmware versions 4.10.xx.xx - 06.15.xx.xx.

Notes:

1. For more information about current Storage Manager controller firmware support for the various DS4000 storage subsystems, see the *IBM DS4000 Storage Manager Installation and User's Guide* for your host operating system.
2. For information about supported host operating systems and operating system requirements, see the Storage Manager readme file for your operating system.

See section “Storage Manager documentation and readme files” on page 1 for instructions that describe how to find the readme files online.

FASTT product renaming

IBM is in the process of renaming some FASTT family products. For a reference guide that identifies each new DS4000 product name with its corresponding FASTT product name, see “FASTT product renaming” on page 2.

Controller firmware 6.14 and 6.15: New features

With controller firmware version **6.15.xx.xx**, the DS4800 storage subsystem utilizes all of the available data cache memory installed in each controller blade for I/O caching. (With controller firmware version 06.14.xx.xx, the DS4800 uses only the first 1 GB of the installed data cache memory per controller for I/O caching.)

Controller firmware versions 6.14.xx.xx and 6.15.xx.xx support:

- DS4800 storage subsystem (only)
- All of the features listed in “Controller firmware 6.10: New features” on page 8 and “Controller firmware 6.12: New features.”

Controller firmware 6.12: New features

Controller firmware **6.12.xx.xx** and later versions all of the features listed in “Controller firmware 6.10: New features” on page 8 in addition to the following new features:

DS4000 FC/SATA Intermix update: Premium Feature Key

DS4000 Storage Manager 9.1x with controller firmware 06.12.xx.xx (and later) supports enabling of the DS4000 FC/SATA Intermix premium feature using a Premium Feature Key.

For more information about using the Intermix premium feature, including configuration and set-up requirements, see the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview*.

New DS4000 Storage Manager installation option

DS4000 Storage Manager 9.1x with controller firmware 06.12.xx.xx (and later) features an installation wizard that enables you to automatically install Storage Manager software packages on your host server.

Note: Using the DS4000 Storage Manager installation wizard requires a system with a graphics card installed. You still have the option of installing the stand-alone host software packages manually. The packages are included with the installation CD.

Support for DS4100 standard (base) SATA Storage Subsystems

Storage Manager 9.1x with controller firmware 06.12.xx.xx (and later) supports DS4100 Standard (Base) SATA Storage Subsystems.

Note: The VolumeCopy, FC/SATA Intermix and Enhanced Remote Mirroring premium features are not supported at this time with the DS4100 Standard (Base) Storage Subsystem. Also, the DS4100 Standard (Base) Storage Subsystem is not supported on AIX host operating systems.

DS4000 Storage Manager usability enhancements

DS4000 Storage Manager 9.12 and later versions feature the following usability enhancements:

- Storage Partitioning wizard, which helps you easily create storage partitions
- Task Assistant, which helps guide you through common enterprise and subsystem management tasks
- Ability to extract SMART data for SATA drives

Controller firmware 6.10: New features

Controller firmware **6.10.xx.xx** and later versions support the following new features:

Enhanced Remote Mirroring

In addition to Metro Mirroring, IBM DS4000 Storage Manager version 9.1x with controller firmware level 6.10.11.xx (and later) also supports Global Copy and Global Mirroring Remote Mirror options. Please see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services User's Guide* for more information.

Note: The terms “Enhanced Remote Mirror Option,” “Metro/Global Remote Mirror Option,” “Remote Mirror,” “Remote Mirror Option,” and “Remote Mirroring” are used interchangeably throughout this document, the SMclient, and the online help system to refer to remote mirroring functionality.

Parallel hard drive firmware download

You can now download drive firmware packages to multiple drives simultaneously, which minimizes downtime. In addition, all files that are associated with a firmware update are now bundled into a single firmware package. See the Subsystem Management window online help for drive firmware download procedures.

Notes:

1. Drive firmware download is an offline management event. You must schedule downtime for the download because I/O to the storage subsystem is not allowed during the drive firmware download process.
2. Parallel hard drive firmware download is not the same thing as concurrent download.

Staged controller firmware download

You can now download the DS4000 controller firmware and NVSRAM to DS4300 Turbo and DS4500 Storage Subsystem for later activation. Depending on your firmware version, DS4000 Storage Subsystem model, and host operating system, the following options might be available:

- Controller firmware download only with immediate activation
- Controller firmware download with the option to activate the firmware at a later time

Note: Staged controller firmware download is not supported on DS4400 Storage Subsystems.

Subsystem Management Window menu enhancements

Troubleshooting, recovery and maintenance tools are now under the Advanced heading in the Subsystem Management window. The following submenus are available:

- Maintenance
- Troubleshooting
- Recovery

Full command-line interface capability

All of the options that are available in SMclient are also available using either the script editor in the Enterprise Management window, or using your preferred command-line interface. For more information about using the command-line interface, see the Enterprise Management window online help.

Support for DS4300 standard (base) Fibre Channel Storage Subsystems

Storage Manager 9.1x with controller firmware 6.10.xx.xx (and later) supports DS4300 Standard (Base) Fibre Channel Storage Subsystems.

Note: The VolumeCopy, FC/SATA Intermix and Enhanced Remote Mirroring premium features are not supported at this time with the DS4300 Standard (Base) Storage Subsystem. Also, the DS4300 Standard (Base) Storage Subsystem is not supported on AIX host operating systems.

DS4000 FC/SATA Intermix premium feature

Storage Manager 9.1x with controller firmware 6.10.xx.xx (and later) supports the DS4000 FC/SATA Intermix premium feature. This premium feature supports the concurrent attachment of Fibre Channel and SATA storage expansion enclosures to a single DS4000 controller configuration.

With controller firmware 6.10.xx.xx and later versions, the FC/SATA Intermix premium feature is enabled using NVSRAM.

For more information about using the Intermix premium feature, including configuration and set-up requirements, see the *IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview* (GC26-7713).

Support for DS4000 EXP710 storage expansion enclosures

Storage Manager 9.1x with controller firmware 6.10.xx.xx (and later) supports DS4000 EXP710 storage expansion enclosures.

Increased support for DS4000 EXP100 SATA storage expansion enclosures

DS4000 EXP100 SATA storage expansion enclosures are now supported on DS4400 Fibre Channel Storage Subsystems.

Also, the DS4100 storage subsystem now supports up to seven EXP100 SATA storage expansion enclosures.

DS4000 Storage Manager usability enhancements

DS4000 Storage Manager 9.10 and later versions feature the following usability enhancements:

- One-click collection of support data, drive diagnostic data, drive channel state management, controller 'service mode,' and the ability to save host topology information
- Improved media error handling for better reporting of unreadable sectors in the DS4000 Storage Subsystem event log, and persistent reporting of unreadable sectors

Storage Manager premium features

The following premium features can be enabled by purchasing a premium feature key. For procedures that describe how to enable premium features, see the *IBM TotalStorage DS4000 Storage Manager Installation and Support Guide* for your operating system, and the Subsystem Management window online help.

Storage Partitioning

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 DS4000 Storage Subsystem models and configurations. For more information about storage partitions, see “Storage partitioning” on page 70.

FlashCopy®

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 61 or see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services Guide*.

VolumeCopy

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 62 or see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services Guide*.

Enhanced Remote Mirroring

The Enhanced Remote Mirroring 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 Enhanced Remote Mirroring option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations. For more information about the Enhanced Remote Mirroring option, see “Enhanced Remote Mirroring option” on page 63 or see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services Guide*.

Fibre Channel/SATA Intermix

The IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix premium feature supports the concurrent attachment of Fibre Channel and SATA storage expansion enclosures to a single DS4000 controller configuration.

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.

Storage subsystem model types

The FAStT500, DS4400, and DS4500, and DS4800 storage subsystem configurations consist of an independent RAID controller enclosure and at least one storage expansion enclosure.

The FAStT200, DS4300 Turbo storage subsystems, and DS4100 and DS4300 base or SCU storage subsystems integrate the storage expansion enclosure and the RAID controller function in the same physical enclosure.

The FAStT200, DS4300 Turbo storage subsystems, DS4800 storage subsystems, and DS4100 and DS4300 base storage subsystems can also be connected with external storage expansion enclosures to increase the storage capacity that can be managed by a single storage subsystem unit.

A DS4000 storage subsystem model might not support the attachment of all available DS4000 drive expansion enclosure models. For example, the DS4800 storage subsystem supports the attachment of the DS4000 EXP710 and EXP100 drive expansion enclosures only. Refer to the *Installation and User's Guide* for your DS4000 storage subsystem model for the supported drive expansion enclosure models for that storage subsystem.

In addition, the DS4000 storage subsystem also support the intermixing of different DS4000 drive expansion enclosure models behind a given DS4000 storage subsystem. There are restrictions, prerequisites, and rules to connect the different drive enclosure models behind a DS4000 storage subsystem. Refer to the *Installation and User's Guide* for your DS4000 storage subsystem model, the DS4000 drive expansion enclosure model and the adding capacity and the *Hard Drive and Storage Expansion Enclosure Installation and Migration Guide* for more information.

The maximum number of drives and storage expansion enclosures that a RAID controller can support depends on the model of the RAID storage subsystems. See the *Installation and User's Guide* for your DS4000 storage subsystem model for the maximum number of drives and storage expansion enclosures that are supported per storage subsystem.

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 storage expansion enclosure	An enclosure that contains drives, power supplies, fans, environmental service modules, assembly, and other supporting components.
Controller	A system board and firmware that control logical drives and implement the storage management functions.
Controller enclosure	An enclosure that contains one or more controllers, power supplies, fans, and other supporting components.

The physical disk capacity of the storage subsystem 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 on page 12 describes the storage

subsystem logical components.

Table 6. Storage subsystem logical components

Component	Description
Array	<p>An array is a set of physical drives that are grouped together logically by the controllers in a storage subsystem. Each array is created with a RAID level to determine how user and redundancy data is written to and retrieved from the drives.</p> <p>The number of drives that can be grouped together into an array depends on the hard drive capacity and the controller firmware version. Each array can be divided into 1 - 256 logical drives.</p>
Logical drive	<p>A logical drive is a logical structure that you create to store data in the DS4000 Storage Subsystem. 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.</p>
Storage partition	<p>A storage partition 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.</p> <p>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 DS4000 Storage Subsystems. See Table 7 on page 13 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 DS4000 Storage Subsystem.</p>
Free capacity	<p>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.</p> <p>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.</p>
Unconfigured capacity	<p>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.</p> <p>Note: In the Subsystem Management window Logical view, unconfigured capacity is displayed as an unconfigured capacity node.</p>

Storage partitioning specifications

Table 7 on page 13 indicates whether the Storage Partition feature is enabled by default and the maximum number of storage partitions that can be defined for a given DS4000 Storage Subsystem.

Table 7. Storage partitioning specifications per DS4000 storage subsystem model

DS4000 subsystem product name (machine type, model number)	Storage Partitioning enabled by default	Maximum number of defined storage partitions	Available storage partition purchase options
DS4800 (1815-82H, 84H, 88H)	Yes (8 partitions standard)	64	8 - 16, 8 - 64, 16 - 64
DS4800 (1815-82A, 84A, 88A)	Choice of 8, 16 or 64	64	8 - 16, 8 - 64, 16 - 64
DS4100 (1724, 100)	No	16	4, 8, 4 - 8, 16, 8 - 16
FAStT200 (3542, 1RU/1RX)	Yes	16	None
DS4300 base (1722, 60U/60X)	No	16	4, 8, 4 - 8, 16, 8 - 16
DS4300 SCU (1724, 6LU/6LX)	No	16	4, 8, 4 - 8, 16, 8 - 16
DS4300 Turbo (1722, 60U/60X)	Yes (8 partitions)	64	8 - 16, 8 - 64, 16 - 64
FAStT500 (3552, 1RU/1RX)	Yes	64	None
DS4400 (1742, 1RU/1RX)	Yes	64	None
DS4500 (1742, 90U/90X)	Yes (16 partitions)	64	16 - 64

Software components

Table 8 describes the IBM TotalStorage DS4000 Storage Manager version 9.1x software components.

Table 8. Storage Manager software components

Component	Description
Storage Manager client (SMclient)	<p>The Storage Manager 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:</p> <ul style="list-style-type: none">• 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. <p>The Storage Manager client is called <i>thin</i> 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.</p> <p>Note: Do not start more than eight instances of the Storage Manager client programs at the same time if the Storage Manager program is installed in multiple host servers or management stations. In addition, do not send more than eight SMcli commands to a storage subsystem at any given time.</p>

Table 8. Storage Manager software components (continued)

Component	Description
Storage Manager host agent (SMagent)	<p>The Storage Manager 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.</p> <p>The host agent, along with the network connection on the host computer, provides an in-band host agent type network-management connection to the storage subsystem instead of the out-of-band direct network-management connection through the individual Ethernet connections on each controller.</p> <p>The management station can communicate with a storage subsystem through the host computer that has host agent management software installed. The host agent receives requests from the management station through the network connection to the host computer, and sends the requests to the controllers in the storage subsystem through the Fibre Channel I/O path.</p> <p>Notes:</p> <ol style="list-style-type: none"> 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. <p>A storage subsystem might be duplicated in the device tree if you are managing it through its Ethernet connections and it is attached to a host computer with the host agent software installed. In this case, you can remove the duplicate storage subsystem icon from the device tree by using the Remove Device option in the Enterprise Management window.</p> Unless you are using Windows NT, you must make a direct (out-of-band) connection to the DS4000 Storage Subsystem in order to set the correct host type. The correct host type will allow the DS4000 Storage Subsystem to configure itself properly for the host server operating system. After you make a direct (out-of-band) connection to the DS4000 Storage Subsystem, depending on your particular site requirements, you can use either or both management methods. Therefore, if you want to manage your subsystem with the in-band management method, you must establish both in-band and out-of-band management connections. <p>Note: Starting with controller firmware 06.14.xx.xx, the default host type is Windows 2000/Server 2003 non-clustered instead of Windows (SP5 or higher) non-clustered.</p>

Table 8. Storage Manager software components (continued)

Component	Description
Redundant disk array controller (RDAC) multipath driver	<p>RDAC is a Fibre Channel I/O path failover driver that is installed on host computers. Usually, a pair of active controllers is located in a storage subsystem enclosure. Each logical drive in the storage subsystem is assigned to a controller. The controller is connected to the Fibre Channel I/O path between the logical drive and the host computer through the Fibre Channel network.</p> <p>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.</p> <p>RDAC requires that the non-failover version of the Fibre Channel Host Bus Adapter is installed in the host server. In addition, the storage subsystem controller must be set to non-ADT mode.</p> <p>Note: The RDAC driver is not available for Hewlett Packard HP-UX and Novell Netware operating systems. In the Novell Netware environment, the Novell native failover driver is used in place of RDAC.</p> <p>Some operating systems, such as HP-UX or VMWARE ESX, have built-in Fibre Channel I/O path failover drivers and do not require this multipath driver. Also, some operating systems support VERITAS® Dynamic Multi-pathing (DMP), which you can choose to use instead of RDAC.</p>

Table 8. Storage Manager software components (continued)

Component	Description
<p>The IBM Storage Area Network (IBMSAN) driver;</p> <ul style="list-style-type: none"> • Netware 5.1 with SP6 or earlier • NetWare 6.0 with SP2 or earlier <p>Netware native failover driver and LSIMPE.cdm for Storage Manager 8.4 and later</p> <ul style="list-style-type: none"> • Netware 5.1 with SP7 or later • NetWare 6.0 with SP3 or later • NetWare 6.5 with SP1.1a and later 	<p>For Netware 5.1 with SP6 or earlier or NetWare 6.0 with SP2 or earlier, 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 DS4000 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 DS4000 disk system controller or expansion unit.</p> <p>The controller ADT function must be enabled for multipath I/O support. To enable the ADT function, set the correct host type for the defined Fibre Channel HBA ports that are installed in the Novell Netware servers. You can do this in the Storage Manager Subsystem Management window. (See “Host types” on page 18 for more information about host types.)</p> <p>Note: IBM currently does not support the IBMSAN as the preferred multipath driver in the Novell Netware operating system environment. You should upgrade the operating system to a later service pack that includes the Novell Netware native multipath driver.</p> <p>For Netware 5.1 with SP7 and later, NetWare 6.0 with SP3 and later or NetWare 6.5 with SP1.1a 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. 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 DS4000 Storage Subsystem or controller enclosure. ADT must be enabled when you use the native Novell Netware multipath driver for LUNs failover and failback.</p> <p>To enable the ADT function, set the correct host type for the defined Fibre Channel HBA ports that are installed in the Novell Netware servers. You can do this in the Storage Manager Subsystem Management window. (See “Host types” on page 18 for more information about host types.)</p> <p>In addition, the LSIMPE.cdm, NWPA.NLM and MM.NLM files must also be installed. Always use the latest version of the LSIMPE.cdm from either the one provided with the IBM DS4000 Fibre Channel HBA device driver or the one that is part of the Novell Netware operating system distribution CD. The LSIMPE.cdm enables the Novell multipath failover driver to identify those logical drives that have been mapped from the DS4000 Storage Subsystem to the host server. Please refer to the Fibre Channel HBA Network driver readme file for more information on how to configure LUN failover and failback.</p>

Table 8. Storage Manager software components (continued)

Component	Description
Storage Manager utility (SMutil)	<p>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:</p> <ul style="list-style-type: none"> • 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. <p>Notes:</p> <ol style="list-style-type: none"> 1. In a Linux operating system environment, you must install RDAC for multi-path failover protection in order to use the utilities in the Storage Manager Utility package. 2. Refer to the Storage Manager readme files for all supported operating systems. See “Storage Manager documentation and readme files” on page 1 for instructions that describe how to find the readme files online.
Linux failover driver, FASTT-MSJ, and QLremote agent (can be used in place of Linux RDAC)	<p>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 DS4000 Storage Subsystem from the controller that was in the failed path to the controller that is in the good path.</p> <p>Note: The controller ADT function must be enabled for multipath I/O support using the Linux failover driver. To enable the ADT function, set the correct host type for the defined Fibre Channel HBA ports that are installed in the Linux servers. You can do this in the Storage Manager Subsystem Management window. See “Host types” for more information.</p>

Host types

The host type setting that you specify when you configure Storage Manager determines how the storage subsystem controllers work with the operating systems on the connected hosts.

All Fibre Channel HBA ports that are defined with the same host type are handled the same way by the DS4000 controllers. This determination is based on the specifications that are defined by the host type. Some of the specifications that differ according to the host type setting include the following options:

Auto Volume Transfer

Enables or disables the Auto-Logical Drive Transfer feature (ADT/AVT). For more information about ADT, see “Auto-Logical Drive Transfer feature” on page 48.

Enable Alternate Controller Reset Propagation

Determines whether the controller will propagate a Host Bus Reset/Target Reset/Logical Unit Reset to the other controller in a dual controller subsystem to support Microsoft Clustering Services.

Allow Reservation on Unowned LUNs

Determines the controller response to Reservation/Release commands that are received for LUNs that are not owned by the controller.

Sector 0 Read Handling for Unowned Volumes

- **Enable Sector 0 Reads for Unowned Volumes:** Applies only to host types with the Auto-Logical Drive Transfer feature enabled. For non-ADT hosts, this option will have no effect.
- **Maximum Sectors Read from Unowned Volumes:** Specifies the maximum allowable sectors (starting from sector 0) that can be read by a controller that does not own the addressed volume. The value of these bits specifies the maximum number of additional sectors that can be read in addition to sector 0.

Reporting of Deferred Errors

Determines how the DS4000 controllers' deferred errors are reported to the host.

Do Not Report Vendor Unique Unit Attention as Check Condition

Determines whether the controller will report a vendor-unique Unit Attention condition as a Check Condition status.

World Wide Name In Standard Inquiry

Enables or disables Extended Standard Inquiry.

Ignore UTM LUN Ownership

Determines how inquiry for the Universal Access LUN (UTM LUN) is reported. The UTM LUN is used by the DS4000 Storage Manager host software to communicate to the DS4000 storage subsystem in DS4000 storage subsystem in-band management configurations.

Report LUN Preferred Path in Standard Inquiry Data

Reports the LUN preferred path in bits 4 and 5 of the Standard Inquiry Data byte 6.

In most DS4000 configurations, the NVSRAM settings for each supported host type for a particular operating system environment are sufficient for connecting a host to the DS4000 storage subsystems. You should not need to change any of the host type settings for NVSRAM. If you think you need to change the NVSRAM settings, please contact your IBM support representative for advice before proceeding.

For information about which host type setting you need to specify for your host operating system and how to specify the setting, see the *IBM Storage Manager Installation and Support Guide* for your operating system.

System requirements

This section provides detailed information about the hardware, software, and storage management architecture for IBM DS4000 Storage Manager Version 9.1x.

Hardware requirements

Table 9 on page 20 lists the hardware that is required to install Storage Manager 9.1x.

Table 9. Storage management architecture hardware components

Component	Description
Management station (one or more)	<p>A management station is a computer that is connected through an Ethernet cable to the host computer or directly to the controller.</p> <ul style="list-style-type: none"> • 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. <p>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.</p>
Network-management station (optional, for SNMP traps)	<p>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).</p> <p>The storage management software sends critical alerts (using SNMP trap messages) to configured destinations.</p>
DHCP/BOOTP or BOOTP-compatible server (for only direct-managed storage subsystems)	<p>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.</p> <p>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.</p>
Host bus adapters (HBAs)	<p>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.</p>
Fibre Channel switches	<p>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.</p> <p>Note: For FASi500, DS4400, and DS4500 Storage Subsystems, 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). This restriction does not apply to the host ports in DS4100, DS4300 and DS4800 Storage Subsystems.</p>
Host computer	<p>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.</p>
Storage subsystem and controller (one or more)	<p>The storage subsystem and storage controller 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).</p>

Table 9. Storage management architecture hardware components (continued)

Component	Description
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.

Storage subsystem management

IBM DS4000 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 DS4000 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.

Note: Do not start more than eight instances of the program at the same time if the DS4000 Storage Manager client program is installed in multiple host server or management stations. You should manage all DS4000 Storage Subsystems in a SAN from a single instance of the Storage Manager client program.

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 1 on page 23.

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 Storage Manager 9.1x.

- 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.

Note: You can avoid DHCP/BOOTP server and network tasks by assigning static IP addresses to the controller, by using a default IP address, or if you are managing all storage subsystems through the Fibre Channel I/O path using a host agent.

Important: Unless you are using Windows NT, you must make direct (out-of-band) connection to the DS4000 Storage Subsystem in order to obtain the correct host type. The correct host type will allow the DS4000 Storage Subsystem to configure itself properly for the host server operating system. After you make a direct (out-of-band) connection to the DS4000 Storage Subsystem, depending on your particular site requirements, you can use either or both management methods. Therefore, if you wish to manage your subsystem with the in-band management method, you must establish both in-band and out-of-band management connections.

To assign a static IP address, if the controller firmware is at 05.3x.xx.xx or earlier, see Retain Tip H171389 “Unable To Setup Networking Without DHCP/BOOTP” at the following Web site:

www.ibm.com/support/

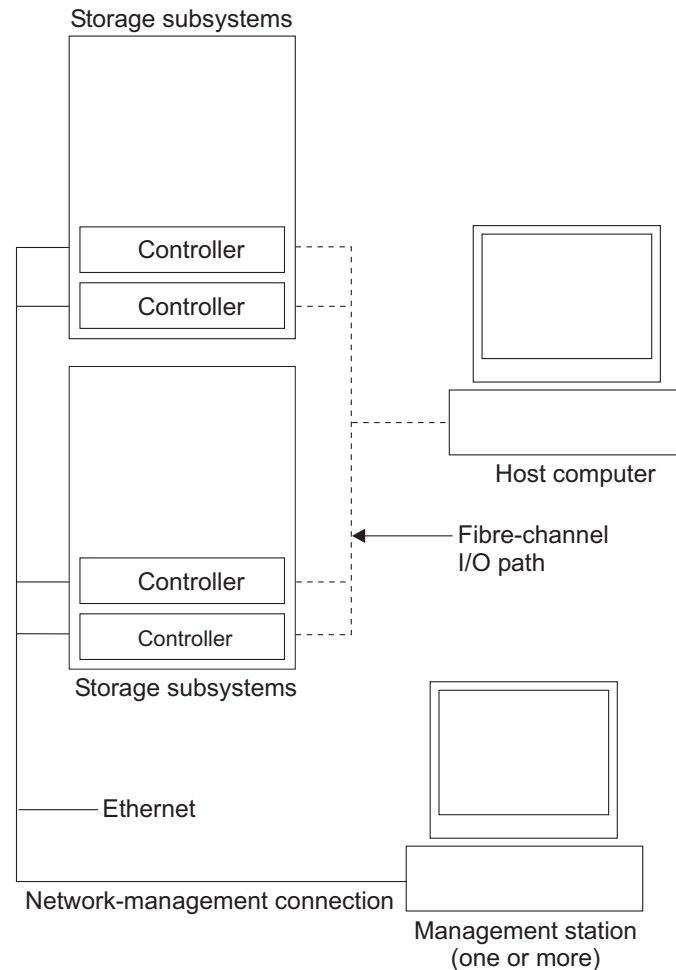
If your controller firmware is at 05.4x.xx.xx or later, you should set the controller static IP address via the SMclient Subsystem Management window after making management connections to the DS4000 controller via in-band or out-of-band management (using the default IP address as indicated in Table 10).

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 (and 192.168.129.101 for DS4800 only)	255.255.255.0
B	192.168.128.102 (and 192.168.129.102 for DS4800 only)	255.255.255.0

Figure 1 shows a system in which the storage subsystems are managed through the direct (out-of-band) management method.



SJ001087

Figure 1. 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.

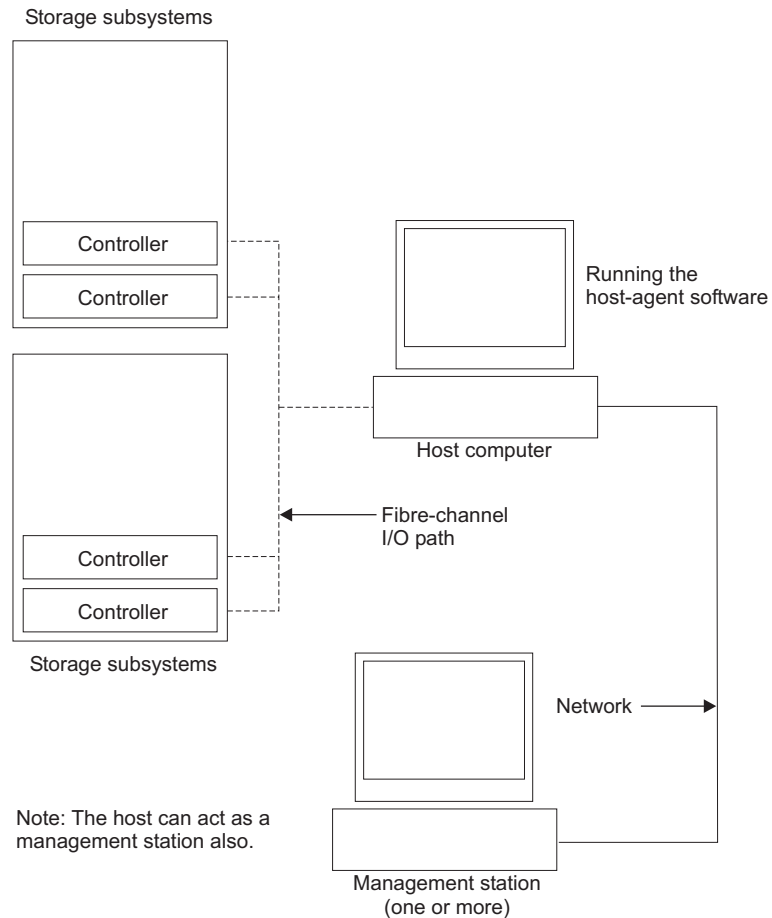
Important:

- 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.
- Systems running the Windows XP operating system can only be used as storage management stations. You cannot use Windows XP as a host operating system.

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 2 on page 25 shows a system in which the storage subsystems are managed through the host agent (in-band) management method.



SJ001089

Figure 2. Host-agent (in-band) managed storage subsystems

Reviewing a sample network

Network A in Figure 3 on page 26 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 3 on page 26 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

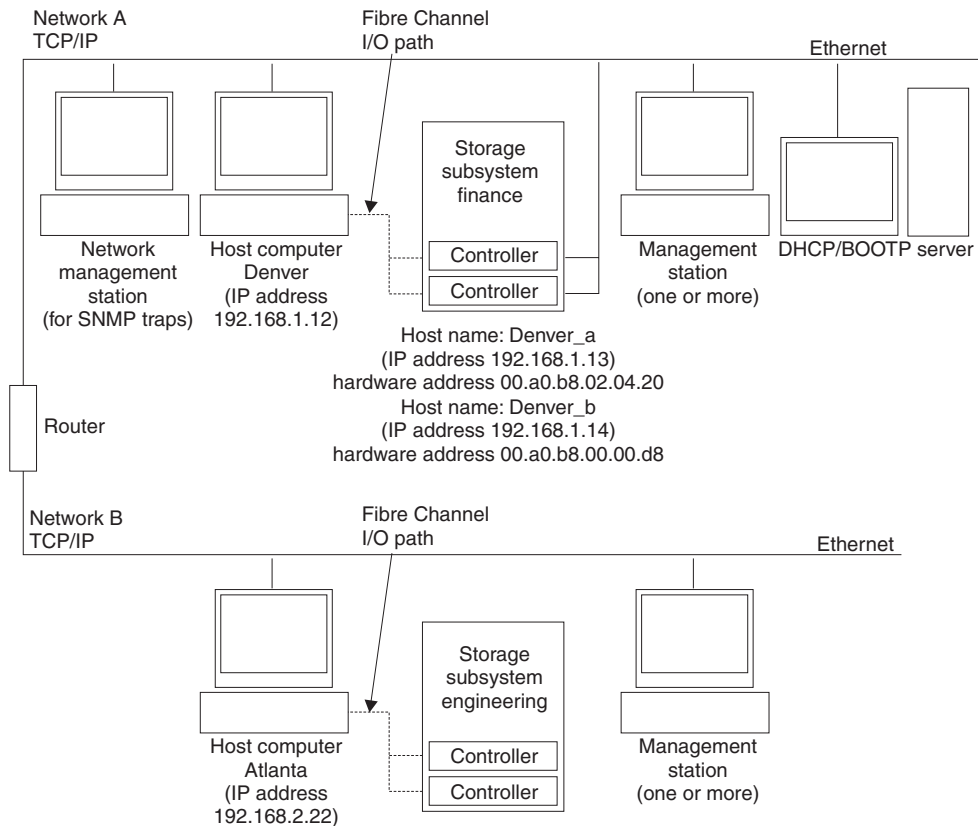


Figure 3. Sample network using direct and host-agent managed storage subsystems

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 06.10.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 7.10
- A storage subsystem with controllers running firmware versions 04.10.xx.xx through 06.1x.xx.xx, which is managed with Storage Manager 9.1x

Important: The common host must have the latest level (version 9.1x) of RDAC and SMagent installed. For DS4300 Turbo, DS4400 and DS4500, the 06.12.xx.xx firmware is available free of charge for download from the IBM support Web site along with all the fixes and software patches. In a coexisting environment, you must upgrade all DS4000 controller firmware to the latest supported code level.

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 4. The Subsystem Management window is shown in Figure 6 on page 30.

Enterprise Management window

The Enterprise Management window, as shown in Figure 4 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
- Configure alert notification destinations, such as email or Simple Network Management Protocol (SNMP) traps, to receive notifications for non-optimal storage subsystems

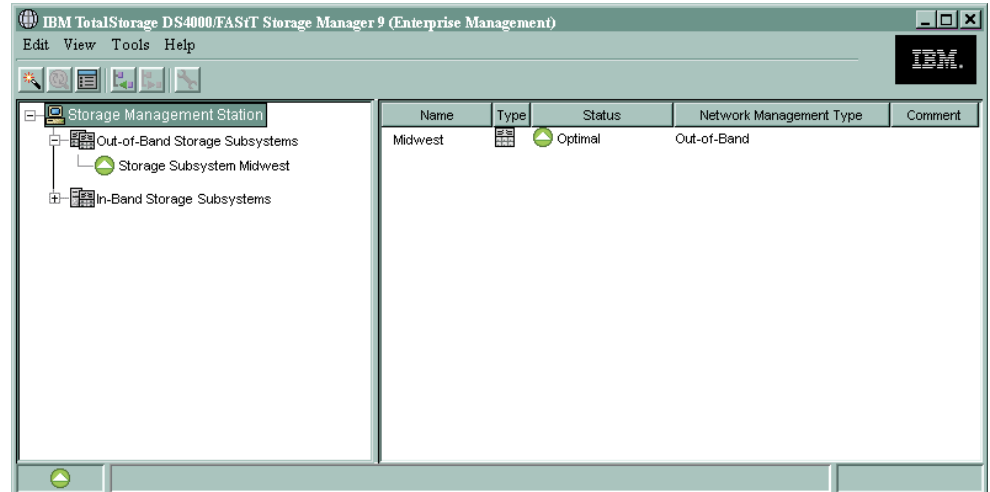


Figure 4. The Enterprise Management window

The emwdata.bin configuration file contains a list of the storage subsystems that are included in the management domain, and any alert destinations you have configured. 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 5.

The Enterprise Management window storage subsystem tree provides a hierarchical view of all the in-band and out-of-band managed storage subsystems. The storage management station node is the root node and sends the storage management commands.

When storage subsystems are added to the Enterprise Management window, they are shown in the device tree as child nodes of the storage management station node. A storage subsystem can be managed through an Ethernet connection on each controller in the storage subsystem (in-band) or through a host interface connection to a host with the host-agent installed (out-of-band).




Figure 5. 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 Storage Subsystem 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 out-of-band management method.

- If a given DS4000 Storage Subsystem is listed in the device tree as being both out-of-band and in-band managed, the DS4000 storage manager program will select the out-of-band route to manage the storage subsystem.

Note: If the DS4000 Storage 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 in the host server that has a Fibre-channel connection to the new storage subsystem. When the host-agent service restarts, the new storage subsystem is detected. Then, go to the Enterprise Management window, select the host server on which you just restarted the host-agent service, and click **Tools** → **Rescan** to add the new storage subsystems to the management domain under the host server node in the device tree.
- 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.
- When storage subsystems are detected or added to the Enterprise Management window for the first time, they are shown as Unnamed in the device tree unless they have been named by another storage management station.

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

Subsystem Management window

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

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

Storage Manager 9.1x supports controller firmware versions 6.1x.xx.xx and 05.xx.xx.xx. To access all of the features of Storage Manager version 9.1x, you must upgrade to the latest controller firmware version that is supported for your DS4000 storagesubsystem model.

Important: Depending on your version of storage management software, the views, menu options, and functionality might differ from the information presented in this guide. For information on available functionality, refer to the documentation supplied with your version of storage management software.

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 subsystem. For example, you can manage two storage subsystems using the Storage Manager software; one storage subsystem has firmware version 6.1x.xx.xx and the other has firmware version 4.xx.xx.xx.

When you open a Subsystem Management window for a particular storage subsystem, the correct Subsystem Management window version is used. The storage subsystem with firmware version 6.1x.xx.xx will use version 9.1x of the storage management software, and the storage subsystem with firmware version 5.30.xx.xx will use version 8.3. 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 maintains the same look and feel of the Storage Manager Subsystem Management window as long as the controller firmware is at a particular version level.

Subsystem Management window tabs

The Subsystem Management window consists of two views: the Logical/Physical view and the Mappings view, as shown in Figure 6. 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.

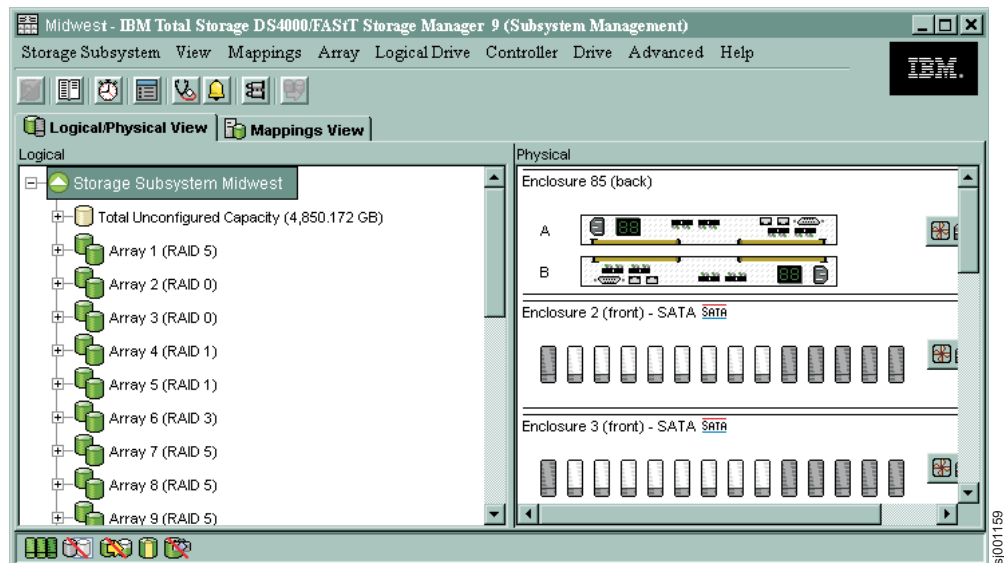


Figure 6. Subsystem Management window Logical View and Physical View

Table 11. Subsystem Management window tabs

Tabs	Description
Logical/Physical View	<p>The Subsystem Management Window Logical/Physical View contains two panes: the Logical View and the Physical View.</p> <p>The Logical View (left pane of Figure 6 on page 30) provides a tree-structured view of logical nodes. This view shows the organization of storage subsystem capacity into arrays and logical drives.</p> <p>The Physical View (right pane of Figure 6 on page 30) provides a view of the physical devices in a storage subsystem, such as controller tray and drive tray components.</p> <p>Selecting a logical drive or other entity in the Logical View shows you the associated physical components in the Physical View.</p> <p>There is a Components button in every controller tray and drive tray that, when clicked, presents the status of each component and shows the temperature status.</p>
Mappings View	<p>The Mappings view of the Subsystem Management window contains two views, Topology and Mappings.</p> <p>The Topology View provides a tree-structured view of logical nodes related to storage partitions.</p> <p>The Mappings View displays the mappings associated with the selected node in the Topology View.</p>

The Subsystem Management window menus

The Subsystem Management window menus are described in Table 12 on page 32. The menus are used to perform storage management operations for a selected storage subsystem or for selected components within a storage subsystem.

Table 12. The Subsystem Management window menus

Menu	Definition
Storage Subsystem	<p>The Storage Subsystem menu contains options to perform the following storage subsystem management operations:</p> <ul style="list-style-type: none"> • Locating functions (locating the storage subsystem by flashing indicator lights) • Automatically configuring the storage subsystem. Save storage subsystem configuration data in a file using the SMcli script commands. • Enabling and disabling premium features • Displaying the Recovery Guru and the corresponding problem summary, details and recovery procedures • Monitoring performance • Changing various Storage Subsystem settings - passwords, default host types, Media scan settings, enclosure order, cache settings and failover alert delay. • Setting controller clocks • Activating or deactivating the Enhanced Remote Mirroring option - Upgrade Mirror Repository Logical Drive • Renaming storage subsystem • Viewing the storage subsystem profile • Managing the controller enclosure alarm (DS4800 only)
View	<p>The View menu allows you to perform the following tasks:</p> <ul style="list-style-type: none"> • Open the Task Assistant tool • Switch the display between the Logical/Physical view and the Mappings view • View associated components for a selected drive in the Physical pane of the Logical/Physical view • Find a particular node in the Logical view or Mappings view • Go directly to a particular Flashcopy, Flashcopy Repository, VolumeCopy source or target logical drive node in the Logical Drive tree.

Table 12. The Subsystem Management window menus (continued)

Menu	Definition
Mappings	<p>The Mappings menu allows you to make changes to or retrieve details about mappings associated with a selected node. The Mappings menu contains the following options:</p> <ul style="list-style-type: none"> • Define hosts, host groups, host ports, or storage partitioning • Change • Move • Replace Host Port • Show All Host Port Information • Remove • Rename <p>Note: You must be in the Mappings View to access the options available in this menu.</p>
Array	<p>The Array menu presents options to perform the following storage management operations on arrays:</p> <ul style="list-style-type: none"> • Locating logical drives • Changing RAID level or controller ownership • Adding free capacity (drives) • Deleting an array <p>Note: These menu options are only available when a an array is selected.</p>
Logical Drive	<p>The Logical Drive menu provides options to perform the following storage management operations on volumes:</p> <ul style="list-style-type: none"> • Creating logical drives • Changing ownership/preferred path, segment size, Media Scan settings, cache settings, modification priority • Increasing capacity • Creating a VolumeCopy • Viewing VolumeCopies using the Copy Manager • Creating, recreating, or disabling a FlashCopy logical drive • Creating, suspending, resuming, or changing remote mirror settings and testing communication. • Removing a mirror relationship • Viewing logical drive properties • Deleting, or renaming a logical drive <p>Note: These menu options are only available when a logical drive is selected.</p>

Table 12. The Subsystem Management window menus (continued)

Menu	Definition
Controller	<p>The Controller menu displays options to perform the following storage management operations on controllers:</p> <ul style="list-style-type: none"> • Changing the preferred loop ID • Modify the IP address, gateway address, or network subnet mask of a controller • Viewing controller properties <p>Note: These menu options are only available when a controller is selected.</p>
Drive	<p>The Drive menu contains options to perform the following storage management operations on drives:</p> <ul style="list-style-type: none"> • Locating a drive and storage expansion enclosure • Assigning or unassigning a hot spare • Viewing drive properties <p>Note: These menu options are only available when a drive is selected.</p>

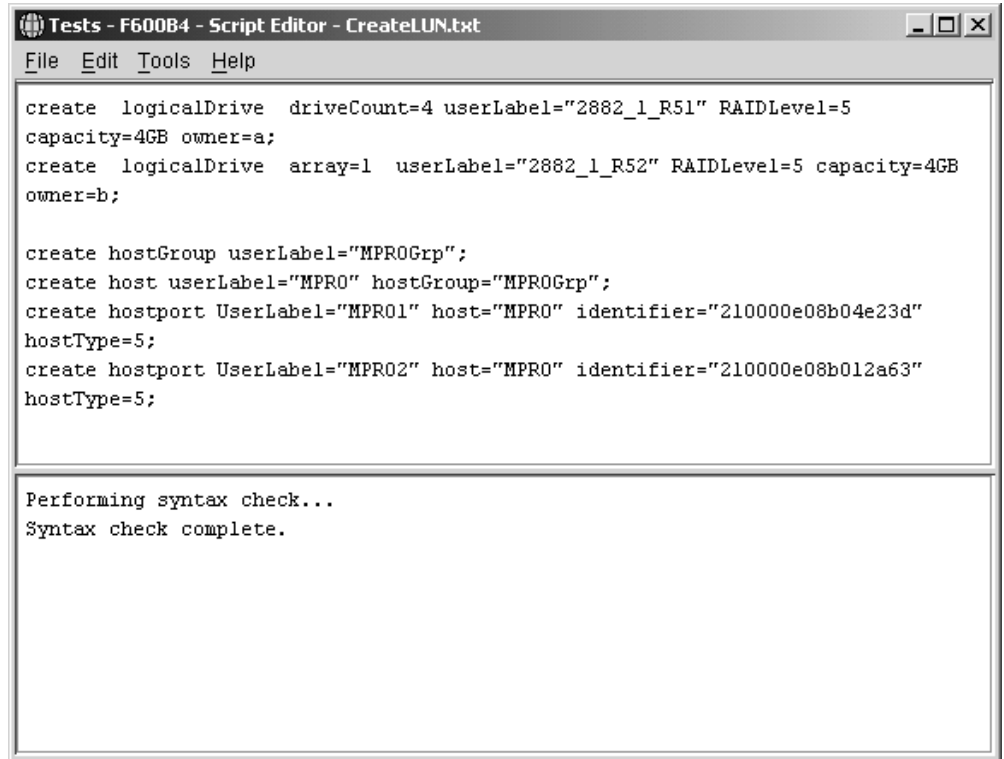
Table 12. The Subsystem Management window menus (continued)

Menu	Definition
Advanced	<p>The Advanced menu allows you to perform certain maintenance functions. The Advanced menu contains the following options:</p> <ul style="list-style-type: none"> • Maintenance <ul style="list-style-type: none"> – Downloading firmware and NVSRAM files – Downloading drive expansion enclosure ESM firmware – Activating or clearing staged controller firmware – Managing persistent reservations – Downloading drive mode pages – Placing an array online or offline • Troubleshooting <ul style="list-style-type: none"> – Collecting support data and drive data – Viewing the event log – Viewing drive channel details – Running Read Link Status diagnostics – Capturing state information – Running controller diagnostics – Running Discrete lines diagnostics (DS4800 only) • Recovery <ul style="list-style-type: none"> – Failing, reconstructing, reviving or initializing a drive – Initializing, revive, and defragment an array – Checking an array for redundancy – Initializing logical drives – Resetting the configuration and controller – Placing controller online, offline, or in service mode – Redistributing logical drives – Displaying unreadable sectors reports – Enabling or disabling data transfer (I/O)
Help	<p>The Help menu provides options to perform the following actions:</p> <ul style="list-style-type: none"> • Display the contents of the Subsystem Management window online 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 7 on page 36, is

provided for running scripted management commands. If the controller firmware version is 5.4x.xx.xx or earlier, some of the management functions that can be done through the GUI are not implemented through script commands. Storage Manager 9.1x in conjunction with controller firmware version 6.10.xx.xx and higher provides full support of all management functions via SMcli commands.



The screenshot shows a window titled "Tests - F600B4 - Script Editor - CreateLUN.txt". The window has a menu bar with "File", "Edit", "Tools", and "Help". The main area contains a script with the following commands:

```
create logicalDrive driveCount=4 userLabel="2882_1_R51" RAIDLevel=5
capacity=4GB owner=a;
create logicalDrive array=1 userLabel="2882_1_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;
```

Below the script, the output shows:

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

SJ001138

Figure 7. 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 controller firmware. The earlier the firmware version, the smaller the set of script commands. For more information about script commands and firmware versions, see the DS4000 Storage Manager Enterprise Management window.

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 editor.

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

In Storage Manager 9.1x with controller firmware 06.10.xx.xx or higher, there is full support for management functions via SMcli commands. For a list of the available commands with the usage syntax and examples, see the Command Reference in the Enterprise window online help.

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 host name or IP address of the controller or controllers
 - For host-agent managed subsystems, enter the host name 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	<p>Add an SNMP trap destination or e-mail alert destination. For example:</p> <p>Use the following command to add an SNMP trap destination:</p> <ul style="list-style-type: none">• <code>-a trap:COMMUNITY,HOST</code> where<ul style="list-style-type: none">– 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. <p>Use the following command to delete an e-mail alert destination:</p> <p><code>-a email:MAILADDRESS</code></p> <p>where MAILADDRESS is the fully qualified e-mail address to which the alert message should be sent</p> <p>Important: There is no space after the colon (:) or the comma (,).</p>

Table 13. Command line parameters (continued)

Command line parameter	Action
-A	<p>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.</p> <p>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.</p>
-c	<p>Specify the list of commands to be performed on the specified storage subsystem.</p> <p>Consider the following usage requirements:</p> <ul style="list-style-type: none"> 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 (;). <p>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.</p>
-d	<p>Display the contents of the configuration file in the following format:</p> <pre><storagearrayname> <hostname> <hostname></pre> <p>The configuration file lists all known storage subsystems that are currently configured in the Enterprise Management window.</p>
-e	Run the commands only, without performing a syntax check first.
-f	<p>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.</p> <p>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.</p>
-F	Specify the e-mail address that will send the alerts.
-i	<p>When used with the -d parameter, display the contents of the configuration file in the following format:</p> <pre><storagearrayname> <IP address> <IP address></pre>
-m	Specify the IP address or host name of the mail or SNMP server that will send the alerts.

Table 13. Command line parameters (continued)

Command line parameter	Action
-n	<p>Specify the storage subsystem name on which you want to perform the script commands.</p> <p>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.</p> <p>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.</p>
-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	<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 circumstances:</p> <ul style="list-style-type: none"> • 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	<p>Delete an SNMP trap destination or e-mail alert destination.</p> <p>Use the following command to delete an SNMP trap destination:</p> <pre>-x trap:COMMUNITY, HOST</pre> <p>where</p> <ul style="list-style-type: none"> • COMMUNITY is the SNMP community name • HOST is the IP address or the host name of a station running an SNMP service. <p>Use the following command to delete an e-mail alert destination:</p> <pre>-x email:MAILADDRESS</pre> <p>where MAILADDRESS is the fully qualified e-mail address to which the alert message should no longer be sent.</p>
-?	Display usage information

Notes:

- All statements must end with a semicolon (;).
- 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 43 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 43.
- 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 host name 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. Storing and 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 Enhanced Remote Mirroring option features are enabled. Standard logical drives are also used with creating FlashCopy logical drives and Enhanced Remote Mirroring 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 61.

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 61.

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 “Enhanced Remote Mirroring option” on page 63.

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 “Enhanced Remote Mirroring option” on page 63.

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 “Enhanced Remote Mirroring option” on page 63.

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 62.

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 62.

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, a minimum of 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.

In addition to these two parameters, you can also specify the segment size, the cache read-ahead count, and which controller is the preferred owner.

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 also 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 DS4000 Storage Subsystem is accomplished with either the Auto-Logical Drive Transfer (ADT) feature, a host multipath driver, or both.

Multipath drivers, such as the redundant disk array controller (RDAC) and VERITAS Volume Manager with Dynamic Multipathing (DMP), are installed on host computers that access the storage subsystem and provide I/O path failover.

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.

Notes:

1. In most cases, ADT is disabled for the operating system for which RDAC is the failover driver. In the “remote boot” configurations, ADT must be enabled.

2. If you are using Dynamic Multi-pathing (DMP) as your default failover driver, you must uninstall RDAC.

Redundant disk array controller (RDAC)

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, Sun Solaris, and Linux (Storage Manager 8.4 and later versions only) 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.

Notes:

1. The RDAC multipath driver is not supported on all operating systems. See the *DS4000 Storage Manager Installation Guide* for your operating system for more information.
2. 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 8.xx and later, this limitation has been relaxed to four paths.
3. If there is more than one path from the host to the storage subsystem, 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 DS4000 Storage Manager software package.

VERITAS Dynamic Multipathing (DMP) is also supported on Windows and Solaris systems.

Novell NetWare

For Storage Manager 9.1x and later, use the Novell native multi-path driver and the LSIMPE.CDM file for failover protection. You must upgrade the installed Netware operating system versions to those that have native failover support because the previous Netware failover solution based on the IBMSAN driver is no longer supported. The following Netware operating system versions have Novell native failover support:

- NW5.1 SP7 and later
- NW 6.0 SP3 and later
- NW 6.5 SP1a and later

With Novell native failover support, 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 DS4000 Storage Subsystem.

For more information on failover protection, see the *IBM TotalStorage DS4000 Storage Manager Installation and User's Guide for Intel-based Operating System Environments* and the readme file of the Fibre-channel HBA Netware device driver.

For the current device driver readme file and failover instructions, go to the following Web site:

www.ibm.com/pc/support

Linux The RDAC driver is provided in the DS4000 Storage Manager software package. It is not supported in all versions of Linux operating system environments. Refer to the RDAC readme for the supported Linux operating system environment for a particular version of Linux RDAC. If you use RDAC as the multipath failover driver, you must ensure that the Fibre Channel HBA device driver is compiled with non-failover setting. In addition, the FAStT-MSJ program, if installed, must be used for diagnostic purposes only.

Instead of using the RDAC driver, you can also use a failover version of the Fibre Channel HBA driver for failover protection. You must use a Fibre Channel diagnostic program called FAStT-MSJ to manually assign logical drives to a preferred path between the DS4000 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.

Linux on POWER-based hosts

for Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server (SLES) on POWER, when using RDAC as a failover driver, the FC HBA device driver must be compiled with non-failover setting. RDAC on Linux is a standalone package located at the following Web site:

www-307.ibm.com/pc/support/site.wss/document.do?Indocid=MIGR-59039

For other POWER-based Linux environments in which RDAC is not supported, you can use a failover version of the Fibre Channel HBA driver for failover protection. Refer to the *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide for AIX, HP-UX, Solaris, and Linux on POWER* for more information on how to set the failover version of the Fibre Channel HBA driver.

AIX and Sun Solaris

The RDAC driver is not included in the DS4000 Storage Manager software package. You must load the "IBM AIX RDAC Driver (fcp.disk.array) separately.

On Solaris, instead of using RDAC, you can use VERITAS DMP.

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 (when using FC HBA failover driver instead of RDAC), 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 that 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 and Linux (when using the RDAC driver and non-failover Fibre-channel HBA driver) 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 re-establish ownership of the logical drive. You must open a storage management window, select Redistribute Logical Drives from the Advanced menu, 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 DS4000 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.

Redundant array of independent disks (RAID)

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 DS4000 controller firmware supports 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 Storage Manager 9.1x software.

Table 14. RAID level configurations

RAID level	Short description	Detailed description
RAID-0	Non-redundant, striping mode	RAID-0 offers simplicity, but does not provide data redundancy. A RAID-0 array spreads data across all drives in the array. This normally provides the best performance but there is not any protection against single drive failure. If one drive in the array fails, all logical drives contained in the array fail. This RAID level is not recommended for high data-availability needs. RAID 0 is better for non-critical data.
RAID-1	Striping/Mirroring mode	<ul style="list-style-type: none"> • A minimum of two drives is required for RAID-1: one for the user data and one for the mirrored data. The DS4000 Storage Subsystem implementation of RAID-1 is basically a combination of RAID-1 and RAID-10, depending on the number of drives selected. If only two drives are selected, RAID-1 is implemented. If you select four or more drives (in multiples of two), RAID 10 is automatically configured across the volume group: two drives for user data, and two drives for the mirrored data. • RAID-1 provides high performance and the best data availability. On a RAID-1 logical drive, data is written to two duplicate disks simultaneously. On a RAID-10 logical drive, data is striped across mirrored pairs. • 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.

Table 14. RAID level configurations (continued)

RAID level	Short description	Detailed description
RAID-3	High-bandwidth mode	<ul style="list-style-type: none"> RAID-3 requires one dedicated disk in the logical drive to hold redundancy information (parity). 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	High I/O mode	<ul style="list-style-type: none"> RAID-5 stripes both user data and redundancy information (parity) across all of the drives in the logical drive. RAID-5 uses the equivalent of one drive's capacity for redundancy information. 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.

To perform a redundancy check, go to **Advanced** → **Recovery** → **Check array redundancy**. 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

Important: A warning box opens when you select the Check array redundancy option that cautions you to only use the option when instructed to do so by the Recovery Guru. It also informs you that if you need to check redundancy for any reason other than recovery, you can enable redundancy checking through Media Scan. For more information on Media Scan, see “Media scan” on page 55.

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. Therefore, 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 for all models of the following DS4000 Storage Subsystems only: FAStT200, FAStT500, DS4100, DS4300, DS4400, and DS4500. There is not any replacement interval for the cache battery backup CRU in other DS4000 Storage Subsystems. 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, DS4100, and DS4300 or DS4300 Turbo disk systems, the battery CRU is located inside each controller CRU. For DS4800, the batteries CRU are located in the Interconnect-batteries 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. Instead, contact IBM service to get a battery replacement as soon as possible to minimize the time that the subsystem is operating with write-caching disabled.

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*.

There are two ways to assign hotspare drives:

Automatically assign drives

If you select this option, hot spare drives are automatically created for the best hot spare coverage using the drives that are available. This option is always available.

Manually assign individual drives

If you select this option, hot spare drives are created out of those drives that were previously selected in the Physical View. This option is not available if you have not selected any drives in the Physical View.

If you choose to manually assign the hotdrives, select a drive with a capacity equal to or larger than the total capacity of the drive you want to cover with the hot spare. For example, if you have an 18 GB drive with configured capacity of 8 GB, you could use a 9 GB or larger drive as a hot spare. Generally, you should not assign a drive as a hot spare unless its capacity is equal to or greater than the capacity of the largest drive on the storage subsystem. For maximum data protection, you should use only the largest capacity drives for hot-spare drives in mixed capacity hard drive configurations.

There is also an option to manually unassign individual drives.

Manually unassign drives

If you select this option, the hot spare drives that you selected in the Physical View are unassigned. This option is not available if you have not selected any drives in the Physical View.

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. Media scan checks the physical disks for defects by reading the raw data from the disk and, if there are errors, writing it back. 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.

Note: 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 DS4000 Storage Subsystem 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.

There are two ways in which media scan can run:

- Background media scan is enabled with logical drive redundancy data checks not enabled.

When redundancy checking is not enabled, the DS4000 Storage Subsystem scans all blocks in the logical drives, including the redundancy blocks, but it does not check for the accuracy of the redundancy data.

This is the default setting when using Storage Manager to create logical drives and it is recommended that you not change this setting.

- Background media scan is enabled with logical drive redundancy data checks enabled.

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.

This setting is not recommended due to the effect redundancy checking has on the server performance.

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

Note: The media scan must be enabled for the entire storage subsystem and enabled on each logical drive within the storage subsystem to protect the logical drive from failure due to media errors.

Media scan only reads data stripes, unless there is a problem. When a block in the stripe cannot be read, the read comment is retried a certain number of times. If the read continues to fail, the controller calculates what that block should be and issues a write-with-verify command on the stripe. As the disk attempts to complete the write command, if the block cannot be written, the drive reallocates sectors until the data can be written. Then the drive reports a successful write and Media Scan checks it with another read. There should not be any additional problems with the stripe. If there are additional problems, the process repeats until there is a successful write, or until the drive is failed due to many consecutive write failures and a hotspare takes over. Repairs are only made on successful writes and the drives are responsible for the repairs. The controller only issues writewithverify commands. Therefore, data stripes can be read repeatedly and report bad sectors but the controller calculates the missing information with RAID.

In a DS4000 dual controller storage subsystem, there are two controllers handling I/O (Controllers A and B). Each logical drive that you create has a preferred controller which normally handles I/O for it. If a controller fails, the I/O for logical drives "owned" by the failed controller fails over to the other controller. Media scan I/O is not impacted by a controller failure and scanning continues on all applicable logical drives when there is only one remaining active controller.

If a drive is failed during the media scan process due to errors, normal reconstruction tasks are initiated in the controllers operating system and Media Scan attempts to rebuild the array using a hotspare drive. While this reconstruction process occurs, no more media scan processing is done on that particular array.

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

- The amount of configured storage capacity in the DS4000 Storage Subsystem

The greater the amount of configured storage capacity in the DS4000 storage subsystem, the higher the performance impact is.

- The configured scan duration for the media scan operations

The longer the scan, the lower the performance impact is.

- The status of the redundancy check option (enabled or disabled)

If redundancy check is enabled, the performance impact is higher due to the need to read the data and recalculated.

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	<p>The drive could not read the data on its first attempt, or on any subsequent attempts.</p> <p>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.</p> <p>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.</p>
Recovered media error	<p>The drive could not read the requested data on its first attempt but succeeded on a subsequent attempt.</p> <p>The data is rewritten to the drive and verified. The error is reported to the event log.</p> <p>Note: Media scan makes three attempts to read the bad blocks.</p>
Redundancy mismatches	<p>Redundancy errors are found.</p> <p>The first 10 redundancy mismatches that are found on a logical drive are reported to the event log.</p> <p>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.</p>
Unfixable error	<p>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.</p> <p>The error is reported to the event log.</p>

Media scan settings

To maximize the protection and minimize the I/O performance impact, the DS4000 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 DS4000 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. Without redundancy check enabled, the controller reads the data stripe to see that all the data can be read. If it reads all the data, it discards the data and moves to the next stripe. When it cannot read a block of data, it reconstructs the data from the remaining blocks and the parity block and issues a write with verify to the block that could not be read. If the block has no data errors, media scan takes the updated information, and verifies that the block was fixed. If the block cannot be rewritten, the drive allocates another block to take the data. When the data is successfully written, the controller verifies that the block is fixed and moves to the next stripe.

Note: With redundancy check, media scan goes through the same process as without redundancy check, but, in addition, the parity block is recalculated and verified. If the parity has data errors, the parity is rewritten. The recalculation and comparison of the parity data requires additional I/O which can affect performance.

Important: Changes to the media settings will not go into effect until the current media scan cycle completes.

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**.

Media scan duration

When media scan is enabled, a duration window is specified (in days) which indicates how long the storage subsystem will give the media scan process to

check all applicable logical drives. The duration window can be shortened or increased to meet the customer requirements. The shorter the duration, the more often a drive is scanned and consequently, the more robust the situation will be. However, the more often a drive is scanned, the higher the performance impact.

Whenever the storage subsystem has some idle time, it starts or continues media scanning operations. If application generated disk I/O work is received, it gets priority. Therefore, the media scan process can slow down, speed up, or in some cases be suspended as the work demands change. If a storage subsystem receives a great deal of application-generated disk I/O, it is possible for the Media Scan to fall behind in its scanning. As the storage subsystem gets closer to the end of the duration window during which it should finish the media scan, the background application starts to increase in priority (i.e. more time is dedicated to the media scan process). This increase in priority only increases to a certain point because the DS4000 Storage Subsystem priority is process application-generated disk I/O. In this case, it is possible that the media scan duration will be longer than the media scan duration settings.

Note: If you change the media scan duration setting, the changes will not take effect until the current media scan cycle completes or the controller is reset.

Copy services and the DS4000 Storage Subsystem

DS4000 Storage Manager 9.1x 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 and higher. 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.

Enhanced Remote Mirroring

The Enhanced Remote Mirroring 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 Enhanced Remote Mirroring option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations. There are two versions of Remote Mirroring Premium features: Remote Mirroring and Enhanced Remote Mirroring (ERM). For more info on ERM, see “Enhanced Remote Mirroring option” on page 63.

Table 16 lists the restrictions that apply to the copy service features.

Table 16. Restrictions to copy services premium feature support

Storage subsystem	Features not supported on controller firmware version 5.3x.xx.xx (See note)	Features not supported on controller firmware version 5.4x.xx.xx	Features not supported on controller firmware version 6.1x.xx.xx
DS4800	N/A	N/A	None Note: DS4800 is supported in the controller firmware 6.1x.xx.xx code thread starting at version 06.14.xx.xx.
DS4100	N/A	Enhanced Remote Mirroring option VolumeCopy	VolumeCopy Note: DS4100 base is supported in the controller firmware 6.1x.xx.xx code thread starting at version 06.12.xx.xx.
DS4100 SCU	N/A	N/A	N/A
DS4300	Enhanced Remote Mirroring option FlashCopy VolumeCopy	Enhanced Remote Mirroring option	Enhanced Remote Mirroring option Note: DS4300 base is supported in the controller firmware 6.1x.xx.xx code thread starting at version 06.12.xx.xx.
DS4300 SCU	Enhanced Remote Mirroring option VolumeCopy	N/A	N/A
DS4300 Turbo	Enhanced Remote Mirroring option VolumeCopy	Enhanced Remote Mirroring option	none Note: DS4300 Turbo is supported in the controller firmware 6.1x.xx.xx code thread starting at version 06.10.xx.xx.
DS4400	VolumeCopy Note: Controller fw 05.3x.xx.xx/05.4x.xx.xx support the first release/version of Remote Mirroring instead of the second version of the RM as supported in controller fw 06.1x.xx.xx.	none	none

Table 16. Restrictions to copy services premium feature support (continued)

Storage subsystem	Features not supported on controller firmware version 5.3x.xx.xx (See note)	Features not supported on controller firmware version 5.4x.xx.xx	Features not supported on controller firmware version 6.1x.xx.xx
DS4500	VolumeCopy Note: Controller fw 05.3x.xx.xx/05.4x.xx.xx support the first release/version of Remote Mirroring instead of the second version of the RM as supported in controller fw 06.1x.xx.xx.	none	none
FAST200	Enhanced Remote Mirroring option VolumeCopy	N/A	N/A
FAST500	VolumeCopy Note: Controller fw 05.3x.xx.xx/05.4x.xx.xx support the first release/version of Remote Mirroring instead of the second version of the RM as supported in controller fw 06.1x.xx.xx.	N/A	N/A

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 DS4000 Storage Manager Version 9 Copy Services User's Guide* or the FlashCopy online help.

VolumeCopy

The VolumeCopy feature is a premium feature that comes with the DS4000 Storage Manager 9.1x 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 DS4000 Storage Manager Version 9 Copy Services User's Guide*.

Enhanced Remote Mirroring option

The following information is an overview of the Enhanced Remote Mirroring option. For more detailed information about the Enhanced Remote Mirroring option, see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services User's Guide*.

The Enhanced Remote Mirroring option is a premium feature that comes with the IBM DS4000 Storage Manager software and is enabled by purchasing a premium feature key. The Enhanced Remote Mirroring option is used for 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 Enhanced Remote Mirroring option enables you to promote a second storage subsystem to take over responsibility for normal I/O operations.

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.

Enhanced Remote Mirroring option enhancements

Many enhancements were made to the remote mirror option with the introduction of Storage Manager 9.1x. The following is a list of these enhancements and a brief description of each. For a more detailed description of the enhancements, see the *IBM TotalStorage DS4000 Storage Manager Version 9 Copy Services User's Guide*.

Delta logging

Delta logging allows the primary array to track the portions of the primary logical drive that have been changed during an inter-array communication interruption.

Suspend and resume

Based on the delta logging framework, a user can manually halt (suspend) mirror synchronization activity to the secondary mirror. The subsequent resume operation attempts to synchronize the data written to the primary logical drive while mirror synchronization was stopped.

Asynchronous write mode

Asynchronous write mode allows the primary-side controller to acknowledge host-initiated write requests before data has been successfully mirrored to the secondaryside controller.

Write order consistency for asynchronous mirrors

For mirror relationships configured for asynchronous write mode, an optional configuration allows the user to specify the mirror to issue write requests to the remote subsystem in the same order as completed on the local subsystem. There are two features that use write order consistency for asynchronous mirrors:

Global Mirroring

If you create multiple remote mirror configurations on your storage subsystem and configure each of the remote mirror pairs to use the asynchronous write mode and preserve write consistency, also known as Global Mirroring, the controller owner treats all of the remote mirror pairs as members of a write consistency group and ensures that the write order is preserved for all remote writes.

Global Copy

If you create multiple remote mirror configurations on your storage subsystem and configure each of the remote mirror pairs to use the asynchronous write mode but do not preserve write consistency, it is known as Global Copy.

Metro Mirroring

The synchronous write mode is now referred to as Metro Mirroring.

Read access to mirror secondary logical drives

This feature allows direct host read access as well as creation of FlashCopy logical drives on mirror secondary logical drives. Read and write access is allowed to FlashCopies of the secondary logical drive.

Enhanced Remote Mirroring diagnostics

There are three new diagnostic services now offered with the Enhanced Remote Mirroring option:

- First, the mirror creation process is improved to provide explicit return status for failed mirror creation requests.
- Second, an inter-subsystem communication diagnostic allows the user to test connectivity between two subsystems after a mirror relationship is in place.
- Third, a new feature also included in this release provides RLS data for host ports. This data can be used to isolate and diagnose intermittent connections at the Fibre Channel level.

Increased number of mirror relationships per subsystem

Storage Manager 9.1x offers 64 mirror relationships per subsystem. However, the increased number of mirrors requires additional logging resources in the mirror repository logical drives. This release creates larger logical drives to accommodate the additional resources, but if smaller repositories exist, the number of mirrors is limited to 32. The user does have the option of expanding his existing repository logical drives so that they can handle 64 volumes.

Resynchronization methods

Two resynchronization methods are available in the current release of the storage management software: Manual Resynchronization, which is the recommended method, and Automatic Resynchronization. Selecting the Manual Resynchronization option allows you to manage the resynchronization process in a way that provides the best opportunity for recovering data.

If a link interruption occurs and prevents communication between the primary logical drive and secondary logical drive in a remote mirror pair, the data on the logical drives might no longer be mirrored correctly. When connectivity is restored between the primary logical drive and secondary logical drive, a resynchronization takes place either automatically or needs to be started manually. During the resynchronization, only the blocks of data that have changed on the primary logical drive during the link interruption are copied to the secondary logical drive.

FlashCopy logical drive enhancement

When creating a FlashCopy in conjunction with Enhanced Remote Mirroring, you are now permitted to base the FlashCopy logical drive on the primary logical drive or secondary logical drive of a remote mirror configuration. This enhancement allows the secondary drive to backup through its FlashCopy image.

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.

Note: There is a limit to how many logical drives you can create in a single storage subsystem. When the Enhanced Remote Mirroring option is enabled, the total number of logical drives that are supported for each storage subsystem is reduced by two from the number of logical drives that you would have without the Enhanced Remote Mirroring 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.

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 write requests.

The new remote mirror option allows the host server to issue read requests to the secondary logical drive.

Note: The host server must have the ability to mount the file system as read-only in order to properly mount and issue read requests to the data in the secondary logical drive.

The secondary logical drive is normally 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 Enhanced Remote Mirroring 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 128 MB for each mirror repository logical drive (256 MB total). If you are upgrading from the previous version of the Enhanced Remote Mirroring option, you must upgrade the size of the repository logical drive from 4 MB to 128 MB in order to support a maximum of 64 remote mirror pairs. Only a maximum of 32 remote mirror pairs is supported with the 4 MB repository logical drive.

Write modes

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 logical drive. The timing of the write I/O completion indication that is sent back to the host depends on the write mode option that is selected.

Asynchronous write mode, which is a new remote mirroring feature, allows the primary-side controller to return the write I/O request completion to the host server before data has been successfully written to the secondary-side controller.

Synchronous write mode, also known as Metro Mirroring, requires that all data has been successfully written to the secondary-side controller before the primary-side controller returns the write I/O request completion to the host server.

Mirror relationships

Before you define a mirror relationship, the Enhanced Remote Mirroring option must be enabled on both the primary and secondary storage subsystems. A secondary standard logical drive candidate (a logical drive that is intended to become one of a mirrored pair) 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.

For more information on the Enhanced Remote Mirroring option, see the *IBM TotalStorage DS4000 Storage Manager Version 9 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.

Important: There is no way to change the password once it is set. Ensure that the password information is kept in a safe and accessible place. Contact IBM technical support for help if you forget the password to the storage subsystem.

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.

Beginning with Storage Manager 9.12 and later versions, in conjunction with controller firmware 06.12 and later, there are Task Wizards in the Enterprise Management and Subsystems Management windows that will guide you through most of the common DS4000 Storage Subsystem management tasks.

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 8.




Figure 8. 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.

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.

For procedures that describe how to create storage partitions and host groups, see the *IBM TotalStorage DS4000 Storage Manager 9 Installation and Support Guide* for your operating system. For more detailed information about storage partitions, see the Subsystem Management window online help.

Switch zoning

You might need to configure switch zoning before you create storage partitions.

Switch zoning is a SAN partitioning method that controls the traffic that runs through a storage networking device, or switch. When you create zones on the switch, the ports outside of a zone are invisible to ports within the zone. In addition, traffic within each zone can be physically isolated from traffic outside the zone.

You can find more information about switch zoning in the *IBM TotalStorage DS4000 Storage Manager Installation and Support Guide* for your operating system.

Storage partitioning terminology

Table 17 describes the storage partitioning terminology that is used in the Mappings view of the Subsystem Management 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	<p>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.</p> <p>The host ports request data from a logical drive on behalf of the host computer; therefore, without associated host ports, a host computer cannot be given a logical drive-to-LUN mapping or request data from a logical drive using a LUN.</p> <p>Initially, all discovered host ports belong to the default host group.</p>
Host computer	<p>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.</p> <p>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.</p>
Host group	<p>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.</p> <p>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.</p>
Default host group	<p>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:</p> <ul style="list-style-type: none"> • Are not involved in specific logical drive-to-LUN mappings • Share access to logical drives with default logical drive-to-LUN mappings
LUN	<p>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.</p>
Specific logical drive-to-LUN mapping	<p>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.</p>

Table 17. Storage partitioning terminology (continued)

Term	Description
Default logical drive-to-LUN mapping	<p>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.</p> <p>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.</p> <p>After installing Storage Manager Version 9.1x software, specific logical drive-to-LUN mappings are created.</p>
Storage partitions mapping preference	<p>You can choose one of the following storage partition mapping preferences when creating a logical drive:</p> <ul style="list-style-type: none"> • Default logical drive-to-LUN mapping • No mapping <p>Choose this option when you create storage partitions to define a specific logical drive-to-LUN mapping for this logical drive.</p>

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.

Obtaining a feature key

Depending on your DS4000 Storage Subsystem 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. For procedures that describe how to enable storage partitioning on your subsystem, see the *IBM TotalStorage DS4000 Storage Manager Installation and Support Guide* for your operating system.

Heterogeneous hosts overview

The Heterogeneous Hosts feature enables host computers that are running different operating systems to access a single storage subsystem.

Note: DS4000 controller firmware versions 04.00.xx.xx and earlier allow 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 2000) or variants of the same operating system (for example, Windows 2000 running in a cluster environment or Windows 2000 running in a noncluster environment). When you specify a host computer type 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.

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 86.
- Recover from storage subsystem problems. See “Recovery Guru” on page 88.

Using the Task Assistant

The Task Assistant provides a convenient, central location from which you can choose to perform the most common tasks in the Enterprise Management window and in the Subsystem Management window.

In the Enterprise Management window, the Task Assistant provides shortcuts to these tasks:

- Adding storage subsystems
- Naming or renaming storage subsystems
- Setting up alert destinations
- Managing storage subsystems

In the Subsystem Management window, the Task Assistant provides shortcuts to these tasks:

- Configuring storage subsystems
- Saving configurations
- Defining hosts
- Creating a new storage partition
- Mapping additional logical drives

If there is a problem with the storage subsystem, a shortcut to the Recovery Guru appears, where you can learn more about the problem and find solutions to correct the problem.

To open the Task Assistant, choose **View >> Task Assistant** from either the Enterprise Management window or the Subsystem Management window, or click the Task Assistant button in the toolbar:



The Task Assistant window opens. See Figure 9 for the Enterprise Management window Task Assistant or Figure 10 on page 77 for the Subsystem Management window Task Assistant.

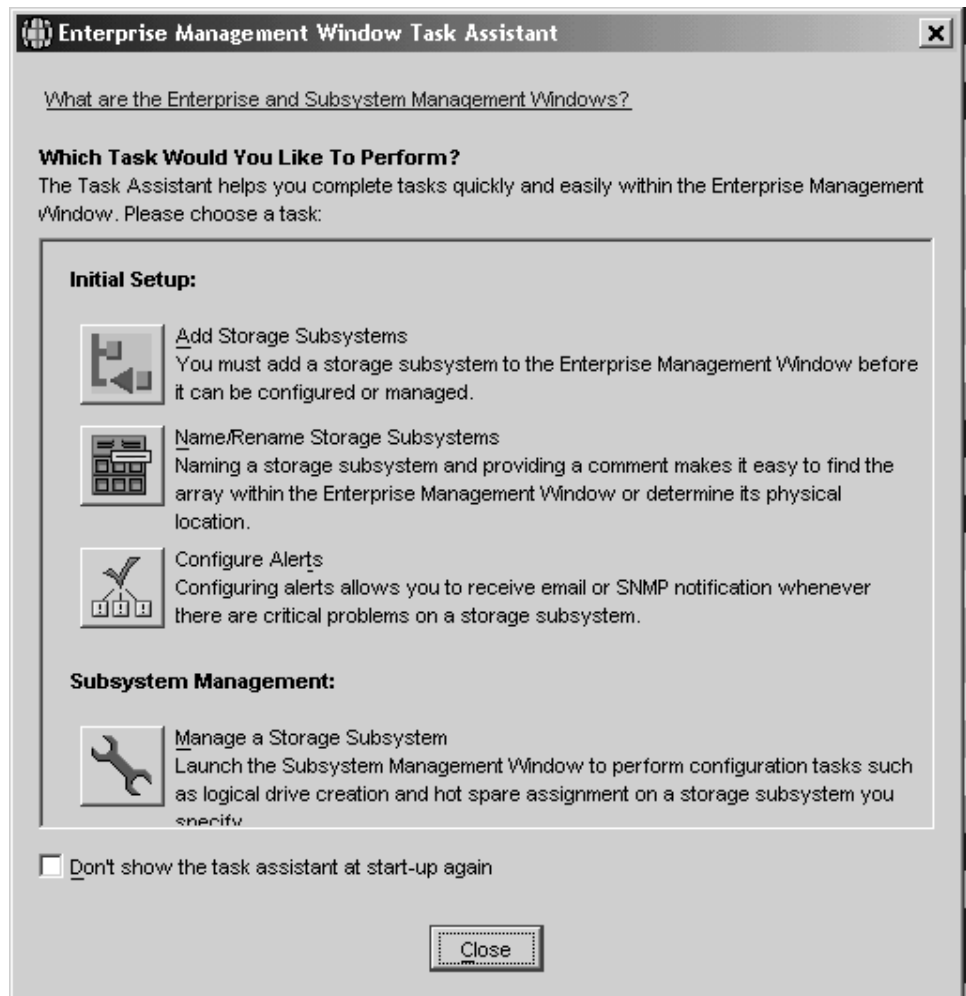


Figure 9. The task assistant in the Enterprise Management window

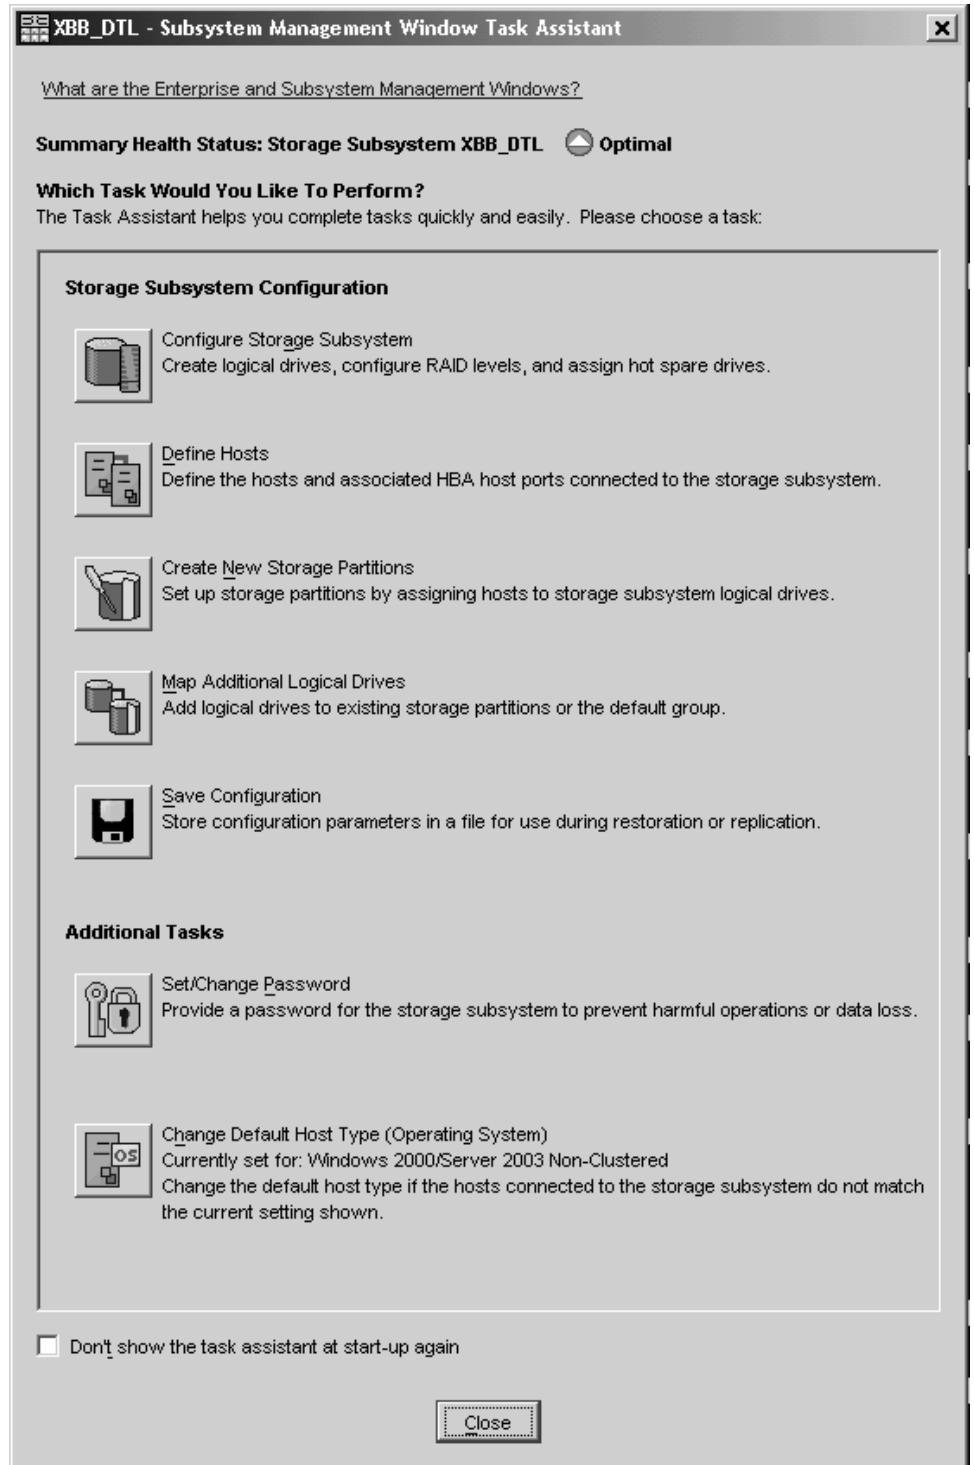


Figure 10. The task assistant in the Subsystem Management window

Note: The Task Assistant is automatically invoked every time you open the Subsystem Management window unless you check the **Don't show the task assistant at start-up again** check box at the bottom of the window.

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 11.

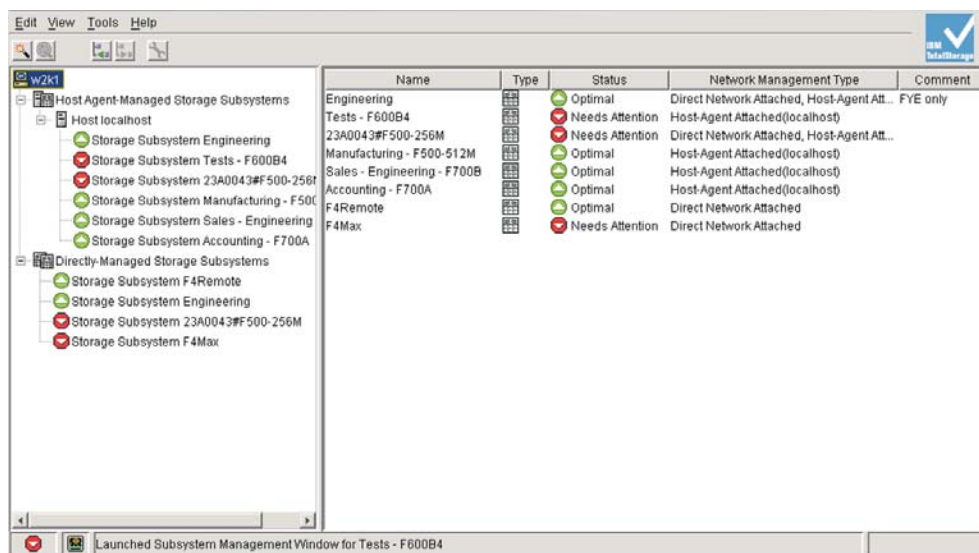


Figure 11. Monitoring storage subsystem health using the Enterprise Management window

Storage subsystem status quick reference



Table 18 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 18. 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).

Table 18. Storage subsystem status icon quick reference (continued)

Icon	Status	Description
	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

The failure notification appears in the Subsystem Management window Logical View.

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

Updating the firmware in the storage subsystem and storage expansion enclosures

The drive firmware, ESM firmware, and controller firmware need to be updated as required. The following sections include information about updating the firmware in the storage subsystem and storage expansion enclosures.

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.

Important: The following sections include information that is useful to know before you download your firmware and NVSRAM. These sections do not include procedures for downloading the firmware and NVSRAM. For detailed instructions on the firmware and NVSRAM downloading procedures, see the *IBM TotalStorage DS4000 Storage Manager Installation and Support Guide* for your operating system.

Attention:

1. IBM supports firmware download with I/O, sometimes referred to as “concurrent firmware download.” Before proceeding with concurrent firmware download, check the readme file packaged with the firmware code or your particular host operating system’s DS4000 Storage Manager host software for any restrictions to this support. See “Storage Manager documentation and readme files” on page 1 for instructions that describe how to find the readme files online.
2. Suspend all I/O activity while downloading firmware and NVSRAM to a DS4000 Storage Subsystem with a single controller or you will not have redundant controller connections between the host server and the DS4000 Storage Subsystem.

Important:

This section provides instructions for downloading DS4000 storage subsystem controller firmware and NVSRAM, DS4000 Storage Expansion Enclosure ESM firmware, and drive firmware. Normally, the DS4000 Storage Subsystem firmware download sequence starts with controller firmware, followed by the NVSRAM and then the ESM firmware, and concludes with the drive firmware. However, always check the DS4000 storage subsystem controller firmware readme file for any controller firmware dependencies and prerequisites before applying the firmware updates to the DS4000 storage subsystem. (See “Storage Manager documentation and readme files” on page 1 for instructions that describe how to find the readme files online.) Updating any components of the DS4000 Storage Subsystem firmware without complying with the dependencies and prerequisites might cause down time (to fix the problems or recover). Recommendation: Contact IBM support if you have any questions regarding the appropriate download sequence for a particular version of firmware.

Downloading controller firmware

There are two methods for downloading the controller firmware, the traditional method and the staged method, which is further described in “The staged controller firmware download feature” on page 81.

Traditional controller firmware download

You must download firmware version 06.1x.xx.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.

Note: The traditional download process takes significantly longer and must be done in one phase, rather than in two phases as with the staged controller firmware download. Therefore the staged controller firmware download, which is described in “The staged controller firmware download feature” on page 81 is the preferred method.

The staged controller firmware download feature

Storage Manager 9.1x, in conjunction with controller firmware version 06.1x.xx.xx or later, offers a new feature in addition to the traditional controller firmware download called the staged controller firmware download.

The staged controller firmware download feature separates firmware loading and firmware activation into two separately executable steps. You can perform the time-consuming task of loading the firmware online so that it is functionally transparent to the application. You can then defer the activation of the loaded firmware to a convenient time. Controller firmware or NVSRAM packages can be downloaded from the storage management software to all storage subsystem controllers. This feature allows you to perform the following actions:

- Controller firmware download only with immediate activation
- NVSRAM download with immediate activation
- Controller firmware download and, optionally, NVSRAM download with the option to activate both later

For more information on the NVSRAM download, see “Downloading NVSRAM.”

Important: Do not perform other storage management tasks, such as creating or deleting logical drives, reconstructing arrays, and so on, while downloading the DS4000 Storage Subsystem controller firmware. It is recommended that you close all storage management sessions (other than the session that you use to upgrade the controller firmware) to the DS4000 Storage Subsystem that you plan to update.

Downloading NVSRAM

There are two methods for downloading the NVSRAM, from a firmware image or from a standalone image. The following sections describe the two methods.

Downloading NVSRAM from a firmware image

If your firmware image contains an NVSRAM image in addition to the firmware executable code, it can be downloaded during the staged controller firmware download. When you attempt to download a firmware image that contains an NVSRAM image onto your storage subsystem, the NVSRAM is programmed into the same flash memory area as the controller firmware executable (i.e. the staging area), and the combined image (of firmware executable and NVSRAM) is treated as a single unit for the purposes of activation. In other words, the NVSRAM image is not copied to the physical NVSRAM until the firmware is activated. After you reboot the controller, both the new executable and the new NVSRAM are active.

Downloading NVSRAM as a standalone image

There is no support for a staged download of an NVSRAM standalone image. When you download a standalone NVSRAM image it is written directly to the physical NVSRAM. An automatic controller reboot is performed after the NVSRAM is written and causes the new NVSRAM to go into effect.

Downloading drive firmware

Important: Do not perform other storage management tasks, such as creating or deleting logical drives, reconstructing arrays, and so on, while downloading the drive firmware. It is recommended that you close all storage management sessions (other than the session that you use to upgrade the firmware) to the DS4000 Storage Subsystem that you plan to update.

General Considerations

The drive firmware download must be done with I/O quiesce. Stop all I/O operations to the storage subsystem before you begin the download process to prevent application errors. Then transfer a downloadable Fibre Channel hard drive firmware file to a drive or drives in the storage subsystem.

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 incompatible 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.

Parallel drive firmware download

The objective of the parallel drive firmware download feature, which is available with Storage Manager 9.1x in conjunction with controller firmware 6.1x.xx.xx or later, is to reduce the data availability impact associated with updating firmware on multiple drives in the storage subsystem. Prior to this feature, the drive firmware download process was issued to one drive at a time. During the drive firmware download cycle, the controllers blocked all I/O access to all logical drives on the subsystem. Updating the drive firmware for all drives in a subsystem could take hours and result in hours of interrupted data availability.

With parallel drive firmware download, a drive firmware image is sent to the controller with a list of drives to update. The controller issues download commands to multiple drives simultaneously. The controller still blocks all I/O access to all logical drives on the subsystem during the download sequence but the overall down time is significantly reduced since multiple drives can be updated concurrently.

A secondary objective of this feature is to simplify the drive firmware download process by bundling all files associated with the firmware update into a single file and providing a mechanism to validate the compatibility of the firmware image with a drive. For example, the typical download sequence for a Fibre Channel drive may require a firmware image and a mode page image. To use the parallel drive firmware download feature, both files are bundled into a single package file.

The following list includes some restrictions and limitations of the parallel drive firmware download feature:

- The maximum number of packages that can be downloaded simultaneously is four.
- The maximum number of drives allowed in one download list is equal to the maximum number of drives that are supported by the storage subsystem.
- A drive cannot be associated with more than one download package in any download command.

- The download of an un-packaged file is not supported.

Environmental services module card

An environmental services module (ESM) card is a customer replaceable unit (CRU) component in the IBM DS4000 storage expansion enclosures (EXP100, EXP500, EXP700, and EXP710) that monitors the environmental condition of the components in that enclosure. It also provides the interface between the Fibre Channel drives in a given storage expansion enclosure with other ESM cards and controller blades in a drive loop.

Downloading ESM firmware

Important: Before you start the ESM firmware download, review the controller firmware readme for any special restrictions for this version of the controller firmware such as certain storage subsystem component hardware or code level requirements. This section does not include procedures for downloading ESM firmware. For downloading procedures, see *IBM TotalStorage DS4000 Storage Manager Installation and Support Guide* for your operating system.

With Storage Manager 9.1x and controller firmware 05.4x.xx.xx or higher, you can update the ESM firmware while host I/O operations are made to the logical drives that are defined in the storage subsystem only if you select to download to one storage expansion enclosure at a time in the ESM firmware download window. The ESM firmware version must be the same in all of the Fibre Channel storage expansion enclosures of the same type in a given DS4000 Storage Subsystem configuration.

For example, if the DS4000 Storage Subsystem has three EXP700 Fibre Channel storage expansion enclosures and two EXP500 Fibre Channel storage expansion enclosures, the firmware of all the ESMs in the two EXP500 Fibre Channel storage expansion enclosures must be the same and the firmware of all ESMs in the three EXP700 Fibre Channel storage expansion enclosures must be the same.

The ESM code for one model (for example, the EXP500) Fibre Channel storage expansion enclosure is not compatible with a different model (for example, EXP700) Fibre Channel storage expansion enclosure.

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

- The IBM Fibre Channel storage expansion enclosures must be connected together in an IBM supported storage expansion enclosure Fibre Channel connection scheme.
- Both of the ESMs in each of the storage expansion 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.

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 storage expansion 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 Global/Metro 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 storage expansion 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 storage expansion enclosure.
- Ensure that the power supplies of the storage expansion 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 12 on page 87) 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 12.

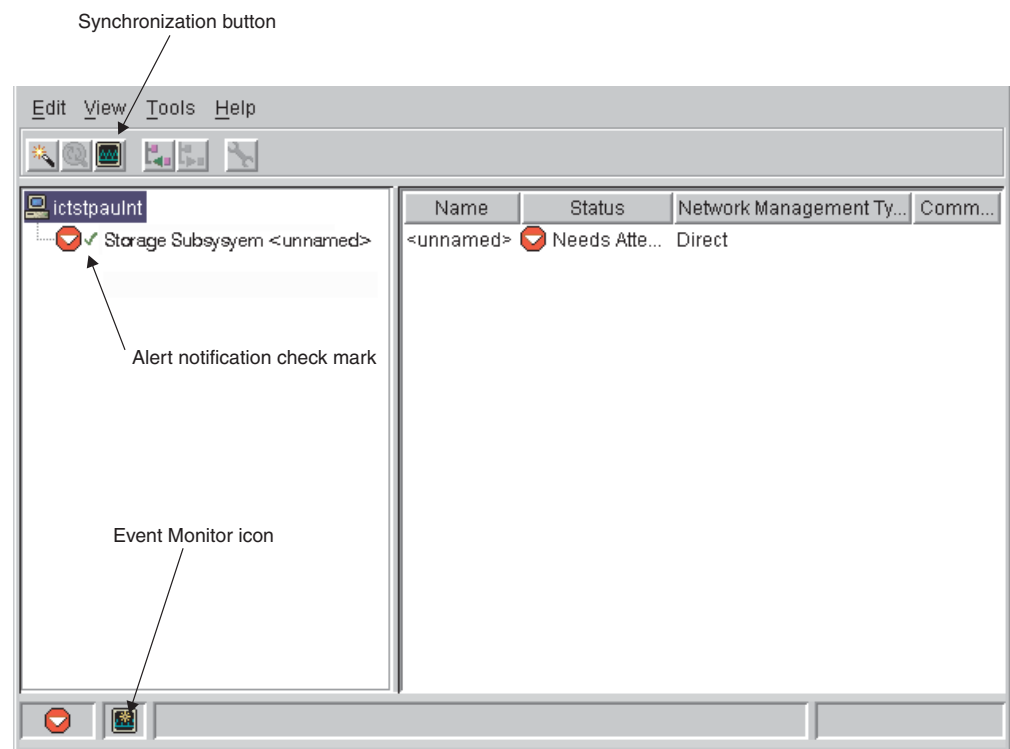


Figure 12. Event monitoring example

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 12 on page 87.

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 13 on page 89, or click **Storage Subsystem** → **Recovery Guru**.

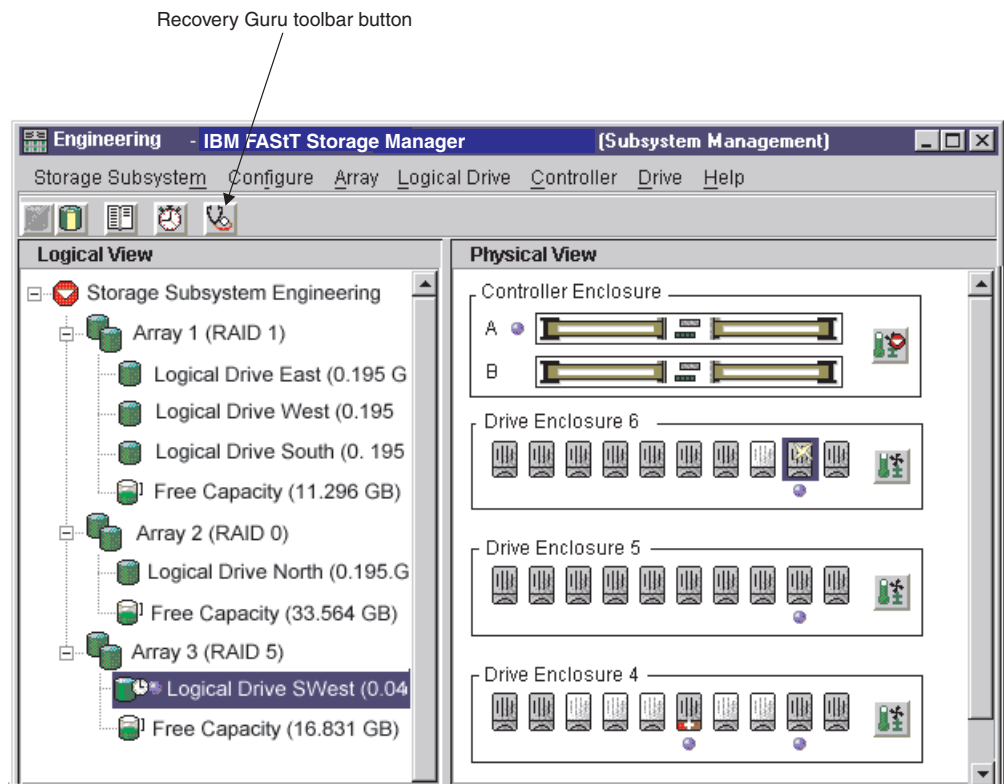


Figure 13. Location of the Recovery Guru toolbar button

The Recovery Guru window is shown in Figure 14. 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.

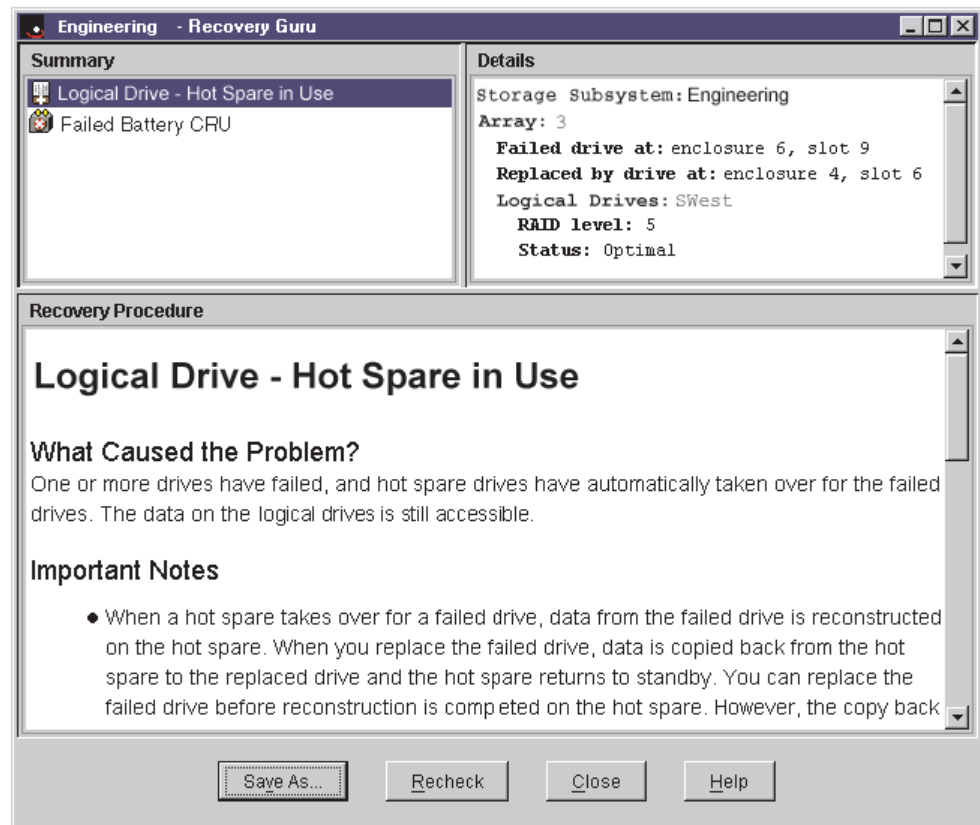


Figure 14. Recovery Guru window

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 15 on page 90.

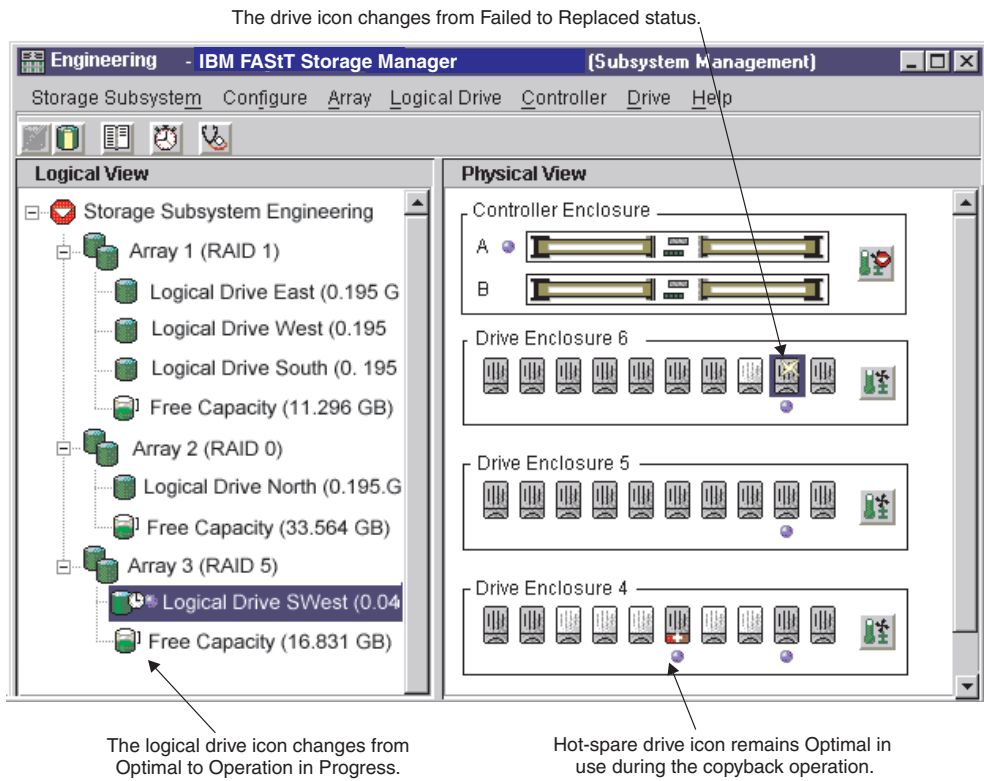
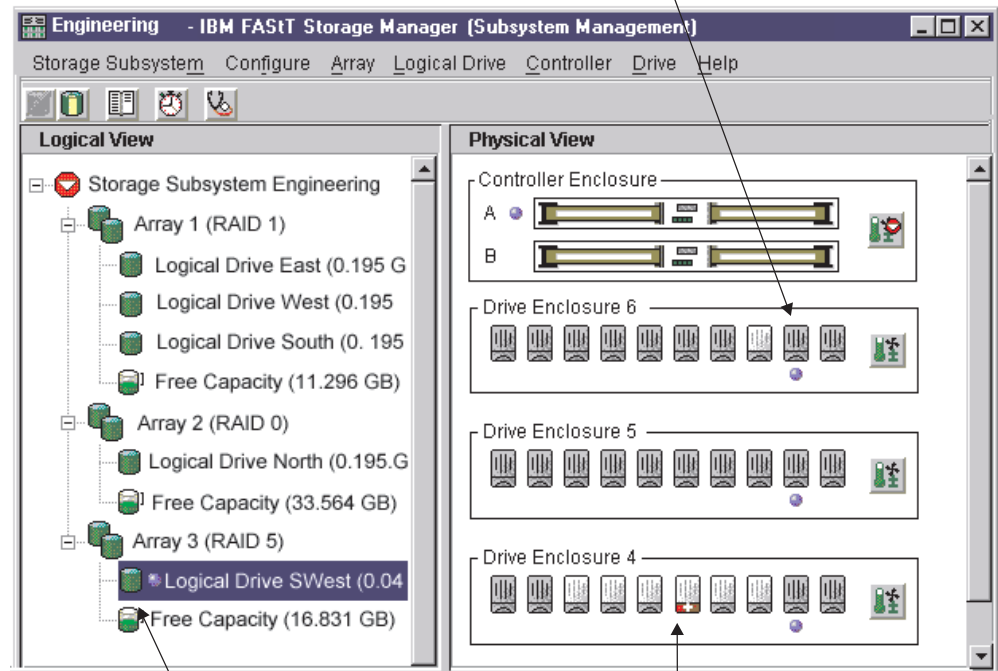


Figure 15. Recovery Guru window showing Replaced status icon

When the copyback operation is complete, the status icon changes to **Optimal**, as shown in Figure 16 on page 91.

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 16. 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 9.1x 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 19 describes the Performance Monitor data that is displayed for selected devices.

Table 19. 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 94.
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 94.
Current KB per second	During the polling interval, the <i>transfer rate</i> is the amount of data, in KB, that is moved through the Fibre Channel I/O path in one second (also called throughput). For more information, see “Optimizing the I/O transfer rate” on page 94.
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 94.
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 94.
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 94.

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 9.1x 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 95. 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 9.1x 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.

If a critical event occurs and you plan to call technical support, you can use the Customer Support Bundle feature to gather and package various pieces of data that can aid in remote troubleshooting. Perform the following steps to use the Customer Support Bundle feature:

1. From the subsystem management window of the logical drive that is exhibiting problems, go to the Advanced menu.
2. Select **Troubleshooting** → **Advanced** → **Collect All Support Data**. The Collect All Support Data window opens.
3. Type the name of the file where you want to save the collected data or click browse to select the file. Click Start.

It takes several seconds for the zip file to be created depending on the amount of data to be collected.

4. Once the process completes, you can send the zip file electronically to customer support for troubleshooting.

Table 20 provides more information about events with a critical priority, as shown in the Subsystem Management window event log.

Table 20. Critical events

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1001 - Channel failed	6/3F/C3	<p>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 relative channel number of the failed channel. Typically this condition is caused by a drive ignoring the SCSI protocol on one of the controller destination channels. The controller fails a channel if it issued a reset on a channel and continues to see the drives ignore the SCSI Bus Reset on this channel.</p> <p>Action: Start the Recovery Guru to access the Failed Drive SCSI Channel recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 1010 - Impending drive failure (PFA) detected	6/5D/80	<p>Description: A drive has reported that a failure prediction threshold has been exceeded. This indicates that the drive might fail within 24 hours.</p> <p>Action: Start the Recovery Guru and click the Impending Drive Failure recovery procedure. Follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1015 - Incorrect mode parameters set on drive	6/3F/BD	<p>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.</p> <p>Action: The controller has not failed yet. Contact your IBM technical-support representative for the instructions to recover from this critical event.</p>
Event 1207 - Fibre-channel link errors - threshold exceeded	None	<p>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.</p> <p>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.</p>
Event 1208 - Data rate negotiation failed	None	<p>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.</p> <p>Action: Start the Recovery Guru to access the Data Rate Negotiation Failed recovery procedure and follow the instructions to correct the failure.</p>
Event 1209 - Drive channel set to Degraded	None	<p>Description: A drive channel status was set to Degraded because of excessive I/O errors or because a technical support representative advised the arrays administrator to manually set the drive channel status for diagnostic or other support reasons.</p> <p>Action: Start the Recovery Guru to access the Degraded Drive Channel recovery procedure and follow the instructions to correct the failure.</p>
Event 150E - Controller loopback diagnostics failed	None	<p>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.</p> <p>Action: Start the Recovery Guru to access the Offline Controller recovery procedure and follow the instructions to replace the controller.</p>
Event 150F - Channel miswire	None	<p>Description: Two or more drive channels are connected to the same Fibre Channel loop. This can cause the storage subsystem to behave unpredictably.</p> <p>Action: Start the Recovery Guru to access the Channel Miswire recovery procedure and follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 1510 - ESM canister miswire	None	<p>Description: Two ESM canisters in the same storage expansion enclosure are connected to the same Fibre Channel loop. A level of redundancy has been lost and the I/O performance for this storage expansion enclosure is reduced.</p> <p>Action: Start the Recovery Guru to access the ESM Canister Miswire recovery procedure and follow the instructions to correct the failure.</p>
Event 1513 - Individual Drive - Degraded Path	None	<p>Description: The specified drive channel is experiencing intermittent errors along the path to a single drive or to several drives.</p> <p>Action: Start the Recovery Guru to access the Individual Drive - Degraded Path recovery procedure and follow the instructions to recover from this failure.</p>
Event 1600 - Uncertified drive detected	None	<p>Description: An uncertified drive has been inserted into the storage subsystem.</p> <p>Action: Start the Recovery Guru to access the Uncertified Drive recovery procedure and follow the instructions to recover from this failure.</p>
Event 1601 - Reserved blocks on ATA drives cannot be discovered	None	<p>Description: Reserved blocks on the ATA drives are not recognized.</p> <p>Action: Contact technical support for instructions on recovering from this event.</p>
Event 200A - Data/parity mismatch detected on logical drive	None	<p>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.</p> <p>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, restore the entire logical drive contents from the most recent backup, if the data is critical.</p>
Event 202E - Read drive error during interrupted write	3/11/8A	<p>Description: A media error has occurred on a read operation during interrupted write processing.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 2109 - Controller cache not enabled - cache sizes do not match	6/A1/00	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 210C - Controller cache battery failed	6/0C/80	<p>Description: The controller has detected that the battery is not physically present, is fully discharged, or has reached its expiration date.</p> <p>Action: Start the Recovery Guru to access the Failed Battery CRU recovery procedure and follow the instructions to correct the failure.</p>
Event 210E - Controller cache memory recovery failed after power cycle or reset	6/0C/81	<p>Description: Recovery from a data-cache error was unsuccessful. User data might have been lost.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2110 - Controller cache memory initialization failed	6/40/81	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2113 - Controller cache battery nearing expiration	6/3F/D9	<p>Description: The cache battery is within six weeks of its expiration.</p> <p>Action: Start the Recovery Guru to access the Battery Nearing Expiration recovery procedure and follow the instructions to correct the failure.</p>
Event 211B - Batteries present but NVSRAM configured for no batteries	None	<p>Description: A battery is present in the storage subsystem but the NVSRAM is set to not include batteries.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2229 - Drive failed by controller	None	<p>Description: The controller failed a drive because of a problem with the drive.</p> <p>Action: Start the Recovery Guru to access the Drive Failed by Controller procedure and follow the instructions to correct the failure.</p>
Event 222D - Drive manually failed	6/3F/87	<p>Description: The drive was manually failed by a user.</p> <p>Action: Start the Recovery Guru to access the Drive Manually Failed procedure and follow the instructions to correct the failure.</p>
Event 2247 - Data lost on the logical drive during unrecovered interrupted write	6/3F/EB	<p>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.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure and follow the instructions to correct the failure. Contact your IBM technical-support representative to complete this procedure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2248 - Drive failed - write failure	6/3F/80	<p>Description: The drive failed during a write command. The drive is marked failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2249 - Drive capacity less than minimum	6/3F/8B	<p>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.</p> <p>Action: Replace the drive with a larger capacity drive.</p>
Event 224A - Drive has wrong block size	6/3F/8C	<p>Description: The drive block size does not match that of the other drives in the logical drive. The drive is marked failed.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224B - Drive failed - initialization failure	6/3F/86	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224D - Drive failed - no response at start of day	6/3F/85	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 224E - Drive failed - initialization/reconstruction failure	6/3F/82	<p>Description: The previously-failed drive is marked failed because of one of the following reasons:</p> <ul style="list-style-type: none"> • 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) <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2250 - Logical drive failure	6/3F/E0	<p>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.</p> <p>Action: Start the Recovery Guru to access the Failed Logical Drive Failure recovery procedure and follow the instructions to correct the failure.</p>
Event 2251 - Drive failed - reconstruction failure	6/3F/8E	<p>Description: A drive failed because of a reconstruction failure during the start-of-day routine.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2252 - Drive marked offline during interrupted write	6/3F/98	<p>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.</p> <p>Action: Start the Recovery Guru to access the Unrecovered Interrupted Write recovery procedure. Contact your IBM technical-support representative to complete this procedure.</p>
Event 2254 - Redundancy (parity) and data mismatch is detected	6/8E/01	<p>Description: The controller detected inconsistent redundancy (parity) or data during a parity verification.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2255 - Logical drive definition incompatible with ALT mode - ALT disabled Note: This event is not applicable for the DS4800.	6/91/3B	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2260 - Uncertified drive	ASC/ASCQ: None	<p>Description: A drive in the storage subsystem is uncertified.</p> <p>Action: Start the Recovery Guru to access the Uncertified Drive recovery procedure.</p>
Event 2602 - Automatic controller firmware synchronization failed	02/04/81	<p>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.</p> <p>Action: Try the firmware download again. If the problem persists, contact your IBM technical-support representative.</p>
Event 2801 - Storage subsystem running on uninterruptible power supply battery	6/3F/C8	<p>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.</p> <p>Action: Start the Recovery Guru and click the Lost AC Power recovery procedure. Follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2803 - Uninterruptible power supply battery - two minutes to failure	6/3F/C9	<p>Description: The uninterruptible power supply has indicated that its standby power supply is nearing depletion.</p> <p>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.</p>
Event 2804 - Uninterruptible power supply battery failed	None	<p>Description: The uninterruptible power supply battery has failed.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 2807 - Environmental service module failed	None	<p>Description: An ESM has failed.</p> <p>Action: Start the Recovery Guru and click the Failed Environmental Service Module CRU recovery procedure. Follow the instructions to correct the failure.</p>
Event 2808 - Enclosure ID not unique	6/98/01	<p>Description: The controller has determined that there are multiple storage expansion enclosures with the same ID selected. Verify that each storage expansion enclosure has a unique ID setting.</p> <p>Action: Start the Recovery Guru and click the Enclosure ID Conflict recovery procedure. Follow the instructions to correct the failure.</p>
Event 280A - Controller enclosure component missing	6/3F/C7	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 280B - Controller enclosure component failed	6/3F/C7	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 280D - Drive expansion enclosures component failed	6/3F/C7	<p>Description: A component other than a drive has failed in the storage expansion enclosure (for example, a fan, power supply, or battery), or an over-temperature condition has occurred. The FRU codes indicate the faulty component.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 280E - Standby power supply not fully charged	6/3F/CA	<p>Description: The uninterruptible power supply has indicated that its standby power supply is not at full capacity.</p> <p>Action: Check the uninterruptible power supply to make sure that the standby power source (battery) is in working condition.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 280F - Environmental service module - loss of communication	6/E0/20	<p>Description: Communication has been lost to one of the dual ESM CRUs in a storage expansion enclosure. The storage expansion enclosure has only one I/O path available.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2813 - Minihub CRU failed	6/3F/C7	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2815 - GBIC failed	None	<p>Description: A gigabit interface converter (GBIC) on either the controller enclosure or the storage expansion enclosure has failed. If there is only one GBIC failure, the storage subsystem is still operational, but a second GBIC failure could result in the failure of the affected enclosures.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2816 - Enclosure ID conflict - duplicate IDs across storage expansion enclosures	6/98/01	<p>Description: Two or more storage expansion enclosures are using the same enclosure identification number.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 2818 - Enclosure ID mismatch - duplicate IDs in the same storage expansion enclosure	6/98/02	<p>Description: A storage expansion enclosure in the storage subsystem contains ESMs with different enclosure identification numbers.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 281B - Nominal temperature exceeded	6/98/03	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 281C- Maximum temperature exceeded	6/3F/C6	<p>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.</p> <p>Action: Start the Recovery Guru and follow the instructions to correct the failure.</p>
Event 281D - Temperature sensor removed	6/98/03	<p>Description: A fan CRU containing a temperature sensor has been removed from the storage subsystem.</p> <p>Action: Replace the CRU as soon as possible. Start the Recovery Guru and click the Failed or Removed Fan CRU recovery procedure and follow the instructions to correct the failure.</p>
Event 281E - Environmental service module firmware mismatch	6/98/03	<p>Description: A storage expansion enclosure in the storage subsystem contains ESMs with different versions of firmware. ESMs in the same storage expansion enclosure must have the same version firmware. If you do not have a replacement service monitor, call your IBM technical-support representative to perform the firmware download.</p> <p>Action: Start the Recovery Guru and click the Environmental Service Module Firmware Version Mismatch recovery procedure. Follow the instructions to correct the failure.</p>
Event 2821 - Incompatible minihub	None	<p>Description: An incompatible minihub canister has been detected in the controller enclosure.</p> <p>Action: Start the Recovery Guru and click the Incompatible Minihub Canister recovery procedure. Follow the instructions to correct the failure.</p>
Event 2823 - Drive bypassed	None	<p>Description: The ESM has reported that the drive has been by-passed to maintain the integrity of the Fibre Channel loop.</p> <p>Action: Start the Recovery Guru to access the By-Passed Drive recovery procedure and follow the instructions to recover from this failure.</p>
Event 2827 - Controller was inadvertently replaced with an ESM	None	<p>Description: A controller canister was inadvertently replaced with an ESM canister.</p> <p>Action: Replace the ESM canister with the controller canister as soon as possible.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2828 - Unsupported storage expansion enclosure selected	None	<p>Description: Your storage subsystem contains one or more unsupported drive enclosures. If all of your drive enclosures are being detected as being unsupported, you might have a problem with an NVSRAM configuration file or you might have the wrong version of firmware. This error condition will cause the drives in the unsupported expansion enclosures to be locked out, which can cause the defined arrays or logical drives to fail.</p> <p>Action: If there are array or logical drive failures, call IBM support for the recovery procedure. Otherwise, Start the Recovery Guru to access the Unsupported Drive Enclosure recovery procedure and follow the instructions to recover from this failure.</p>
Event 2829 - Controller redundancy lost	6/E0/20	<p>Description: Communication has been lost between the two controllers through one of the drive loops (channels).</p> <p>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.</p>
Event 282B - storage expansion enclosure path redundancy lost	6/E0/20	<p>Description: A storage expansion 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.</p> <p>Action: Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.</p>
Event 282D - Drive path redundancy lost	6/E0/20	<p>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.</p> <p>Action: Start the Recovery Guru and click the Drive - Loss of Path Redundancy recovery procedure. Follow the instructions to correct the failure.</p>
Event 282F - Incompatible version of ESM firmware detected	None	<p>Description: A storage expansion enclosure in the storage subsystem contains ESM canisters with different firmware versions. This error might also be reported if a storage expansion enclosure in the storage subsystem contains ESM canisters with different hardware.</p> <p>Action: Start the Recovery Guru to access the ESM Canister Firmware Version Mismatch recovery procedure and follow the instructions to recover from this failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 2830 - Mixed drive types not supported	None	<p>Description: The storage subsystem currently contains drives of different drive technologies, such as Fibre Channel (FC) and Serial ATA (SATA). Mixing different drive technologies is not supported on this storage subsystem.</p> <p>Action: Select the Recovery Guru to access the Mixed Drive Types Not Supported recovery procedure and follow the instructions to recover from this failure.</p>
Event 2835 - Drive expansion enclosures not cabled together	ASC/ASCQ: None	<p>Description: There are drive expansion enclosures in the storage subsystem that are not cabled correctly; they have ESM canisters that must be cabled together sequentially.</p> <p>Action: Start the Recovery Guru to access the Drive Enclosures Not Cabled Together recovery procedure and follow the instructions to recover from this failure.</p>
Event 3019 - Logical drive ownership changed due to failover	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.</p>
Event 4011 - Logical drive not on preferred path	None	<p>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).</p> <p>Action: Start the Recovery Guru and click the Logical Drive Not on Preferred Path recovery procedure. Follow the instructions to correct the failure.</p>
Event 5005 - Place controller offline	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 502F - Missing logical drive deleted	None	<p>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 storage expansion enclosures.</p> <p>Action: Start the Recovery Guru and click the Missing Logical Drive recovery procedure. Follow the instructions to correct the failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 5038 - Controller in lockout mode	None	<p>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.</p> <p>Action: Wait 10 minutes and try to enter the password again.</p>
Event 5040 - Place controller in service mode	None	<p>Description: The controller was manually placed in service mode for diagnostic or recovery reasons.</p> <p>Action: Start the Recovery Guru to access the Controller in Service Mode recovery procedure. Use this procedure to place the controller back online.</p>
Event 5405 - Gold Key - mismatched settings	ASC/ASCQ: None	<p>Description: Each controller in the controller pair has a different NVSRAM bit setting that determines if the controller is subject to Gold Key restrictions.</p> <p>Action: This critical event should not be seen in the IBM DS4000 Storage Subsystem configuration. This event could be generated if there is an inadvertent swapping of IBM storage subsystem controllers or drives with non-IBM controllers or drives. Contact IBM Support for the recovery procedure.</p>
Event 5406 - Mixed drive types - mismatched settings	ASC/ASCQ: None	<p>Description: Each controller in the controller pair has a different setting for the NVSRAM bit that controls whether Mixed Drive Types is a premium feature.</p> <p>Action: Start the Recovery Guru to access the Mixed Drive Types - Mismatched Settings recovery procedure and follow the instructions to correct this controller condition.</p>
Event 5602 - This controller's alternate failed - timeout waiting for results	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 560B - CtlrDiag task cannot obtain Mode Select lock	None	<p>Description: This controller is attempting to run diagnostics and could not secure the test area from other storage subsystem operations. The diagnostics were canceled.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 560C - CtlrDiag task on controller's alternate cannot obtain Mode	None	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 560D - Diagnostics read test failed on controller	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 560E - This controller's alternate failed diagnostics read test	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 560F - Diagnostics write test failed on controller	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 5610 - This controller's alternate failed diagnostics write test	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Offline Controller recovery procedure. Follow the instructions to replace the controller.</p>
Event 5616 - Diagnostics rejected - configuration error on controller	None	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 5617 - Diagnostics rejected - configuration error on controller's alternate	None	<p>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.</p> <p>Action: Contact your IBM technical-support representative for the instructions to recover from this failure.</p>
Event 6101 - Internal configuration database full	None	<p>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:</p> <ul style="list-style-type: none"> • FlashCopy logical drive configuration data • Global/Metro remote mirror configuration data <p>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.</p>
Event 6107 - The alternate for the controller is non-functional and is being held in reset	None	<p>Description: A controller in the storage subsystem has detected that its alternate controller is non-functional due to hardware problems and needs to be replaced.</p> <p>Action: Start the Recovery Guru to access the Offline Controller recovery procedure and follow the instructions to recover from this failure.</p>
Event 6200 - FlashCopy repository logical drive threshold exceeded	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the FlashCopy Repository Logical Drive Threshold Exceeded recovery procedure. Follow the instructions to correct this failure.</p>
Event 6201 - FlashCopy repository logical drive full	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the FlashCopy Repository Logical Drive Capacity - Full recovery procedure. Follow the instructions to correct this failure.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6202 - Failed FlashCopy logical drive	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Failed FlashCopy Logical Drive recovery procedure. Follow the instructions to correct this failure.</p>
Event 6400 - Dual primary logical drive	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Dual Primary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.</p>
Event 6401 - Dual secondary logical drive	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Dual Secondary Logical Drive Conflict recovery procedure. Follow the instructions to correct this failure.</p>
Event 6402 - Mirror data unsynchronized	Not recorded with event	<p>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.</p> <p>Action: Start the Recovery Guru and click the Mirror Data Unsynchronized recovery procedure. Follow the instructions to correct this failure.</p>
Event 6503 - Remote logical drive link down	None	<p>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.</p> <p>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.</p>

Table 20. Critical events (continued)

Critical event number	Sense key/ASC/ASCQ	Critical event description and required action
Event 6505 - WWN change failed	None	<p>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.</p> <p>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.</p>
Event 6600 - Logical drive copy operation failed	None	<p>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.</p> <p>Action: Start the Recovery Guru and click the Logical Drive Copy Operation Failed recovery procedure. Follow the instructions to correct this failure.</p>
Event 6700 - Unreadable sector(s) detected - data loss occurred	None	<p>Description: Unreadable sectors have been detected on one or more logical drives and data loss has occurred.</p> <p>Action: Start the Recovery Guru to access the Unreadable Sectors Detected recovery procedure and follow the instructions to recover from this failure.</p>
Event 6703 - Overflow in unreadable sector database	None	<p>Description: The Unreadable Sectors log has been filled to its maximum capacity.</p> <p>Action: Select the Recovery Guru to access the Unreadable Sectors Log Full recovery procedure and follow the instructions to recover from this failure.</p>

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)
	Changing logical drive and array properties (name, segment size, cache settings, media scan settings, preferred/owner controller)
	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
Creating a VolumeCopy target logical drive. Managing VolumeCopy logical drive pair	
Activating/Deactivating Enhanced Remote Mirroring feature. Creating and managing Enhanced Remote Mirroring logical drive pairs	
Creating FlashCopy logical drive(s). Managing existing FlashCopy logical drives (deleting Flashcopy logical drive, re-creating Flashcopy logical drive or increasing Flashcopy Repository logical drive capacity)	

Using the Script Editor

In Storage Manager version 9.1x with controller firmware 06.1x.xx.xx, there is full support for management functions via SMcli commands. For a list of the available commands with the usage syntax and examples, see the Command Reference in the Enterprise window online help.

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
	Undefined 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 Running and displaying Drive Channel diagnostics Capturing all support data and storage subsystem state information

Recovering from problems

If a critical event occurs and you plan to call technical support, you can use the Customer Support Bundle feature to gather and package various pieces of data that can aid in remote troubleshooting. For more information on the Customer Support Bundle feature, see 97.

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
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 DS4000 documentation

The following tables present an overview of the DS4500, DS4400, DS4300, DS4100, and DS4800 Storage Subsystem product libraries, as well as other related documents. Each table lists documents that are included in the libraries and what common tasks they address.

You can access the documents listed in these tables at one of the following Web sites:

www.ibm.com/servers/storage/support/disk/

www.ibm.com/shop/publications/order/

DS4500 Fibre Channel Storage Server library

Table 21 associates each document in the DS4500 (previously FAStT900) Fibre Channel Storage Server library with its related common user tasks.

Table 21. TotalStorage DS4500 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM TotalStorage DS4500 Installation and Support Guide, GC26-7727</i>	✓	✓		✓		
<i>IBM TotalStorage DS4500 Fibre Channel Cabling Instructions, GC26-7729</i>	✓	✓				
<i>IBM TotalStorage DS4500 Storage Server User's Guide, GC26-7726</i>				✓	✓	✓
<i>IBM TotalStorage DS4500 Rack Mounting Instructions, GC26-7728</i>	✓	✓				

DS4400 Fibre Channel Storage Server library

Table 22 associates each document in the DS4400 (previously FAStT700) Fibre Channel Storage Server library with its related common user tasks.

Table 22. TotalStorage DS4400 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM DS4400 Fibre Channel Storage Server User's Guide, GC26-7730</i>	✓	✓		✓	✓	✓
<i>IBM DS4400 Fibre Channel Storage Server Installation and Support Guide, GC26-7731</i>	✓	✓		✓	✓	
<i>IBM DS4400 Fibre Channel Cabling Instructions, GC26-7732</i>	✓	✓				

DS4300 Fibre Channel Storage Server library

Table 23 associates each document in the DS4300 (previously FAStT600) Fibre Channel Storage Server library with its related common user tasks.

Table 23. TotalStorage DS4300 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM TotalStorage DS4300 Fibre Channel Storage Server Installation and User's Guide, GC26-7722</i>	✓	✓		✓		
<i>IBM TotalStorage DS4300 Rack Mounting Instructions, GC26-7724</i>	✓	✓				
<i>IBM TotalStorage DS4300 Fibre Channel Cabling Instructions, GC26-7725</i>	✓	✓				
<i>IBM TotalStorage DS4300 SCU Base Upgrade Kit, GC26-7740</i>		✓	✓			
<i>IBM TotalStorage DS4300 SCU Turbo Upgrade Kit, GC26-7741</i>		✓	✓			
<i>IBM TotalStorage DS4300 Turbo Models 6LU/6LX Upgrade Kit, GC26-7723</i>		✓	✓			

DS4100 SATA Storage Server library

Table 24 associates each document in the DS4100 (previously FAStT100) SATA Storage Server library with its related common user tasks.

Table 24. TotalStorage DS4100 SATA Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM TotalStorage DS4100 Installation, User's and Maintenance Guide, GC26-7733</i>	✓	✓		✓	✓	✓
<i>IBM TotalStorage DS4100 Cabling Guide, 24P8973</i>		✓				

DS4800 Storage Subsystem library

Table 25 associates each document in the DS4800 Storage Subsystem library with its related common user tasks.

Table 25. TotalStorage DS4800 Storage Subsystem document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM TotalStorage DS4800 Installation, User's and Maintenance Guide, GC26-7748</i>	✓	✓		✓	✓	✓
<i>IBM TotalStorage DS4800 Installation and Cabling Overview, 22R4738</i>		✓				

Other DS4000 and DS4000-related documents

Table 26 associates each of the following documents with its related common user tasks.

Table 26. TotalStorage DS4000 and DS4000-related document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM Safety Information, 33P2406</i>					✓	
<i>IBM TotalStorage DS4000 Quick Start Guide, GC26-7738</i>	✓	✓				
<i>IBM TotalStorage DS4000 Hardware Maintenance Manual, GC26-7702</i>						✓
<i>IBM TotalStorage DS4000 Problem Determination Guide, GC26-7703</i>						✓
<i>IBM Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329</i>	✓	✓			✓	✓
<i>IBM TotalStorage DS4000 FC2-133 Host Bus Adapter Installation and User's Guide, GC26-7736</i>		✓			✓	
<i>IBM TotalStorage DS4000 FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7737</i>		✓			✓	
<i>IBM TotalStorage DS4000 Fibre Channel and Serial ATA Intermix Premium Feature Installation Overview GC26-7713</i>	✓	✓	✓	✓		
<i>Fibre Channel Solutions - IBM DS4000 EXP500 Installation and User's Guide, 59P5637</i>	✓	✓		✓	✓	✓

Table 26. TotalStorage DS4000 and DS4000–related document titles by user tasks (continued)

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
<i>IBM TotalStorage DS4000 EXP700 and EXP710 Storage Expansion Enclosures Installation, User's, and Maintenance Guide, GC26-7735</i>	✓	✓		✓	✓	✓
<i>IBM TotalStorage DS4000 Hard Drive and Storage Expansion Enclosures Installation and Migration Guide, GC26-7704</i>	✓	✓				
<i>IBM DS4000 Management Suite Java User's Guide, 32P0081</i>					✓	✓
<i>IBM Netfinity Fibre Channel Cabling Instructions, 19K0906</i>		✓				
<i>IBM Fibre Channel SAN Configuration Setup Guide, 25P2509</i>	✓		✓	✓	✓	

Appendix C. Accessibility

This section provides information about alternate keyboard navigation, which is a DS4000 Storage Manager accessibility feature. Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

By using the alternate keyboard operations that are described in this section, you can use keys or key combinations to perform Storage Manager tasks and initiate many menu actions that can also be done with a mouse.

Note: In addition to the keyboard operations that are described in this section, the DS4000 Storage Manager 9.14 software installation package for Windows includes a screen reader software interface. To enable the screen reader, select Custom Installation when using the installation wizard to install Storage Manager 9.14 on a Windows host/management station. Then, in the Select Product Features window, select Java Access Bridge in addition to the other required host software components.

Table 27 defines the keyboard operations that enable you to navigate, select, or activate user interface components. The following terms are used in the table:

- *Navigate* means to move the input focus from one user interface component to another.
- *Select* means to choose one or more components, typically for a subsequent action.
- *Activate* means to carry out the action of a particular component.

Note: In general, navigation between components requires the following keys:

- **Tab** - Moves keyboard focus to the next component or to the first member of the next group of components
- **Shift-Tab** - Moves keyboard focus to the previous component or to the first component in the previous group of components
- **Arrow keys** - Move keyboard focus within the individual components of a group of components

Table 27. DS4000 Storage Manager alternate keyboard operations

Short cut	Action
F1	Open the Help.
F10	Move keyboard focus to main menu bar and post first menu; use the arrow keys to navigate through the available options.
Alt+F4	Close the management window.
Alt+F6	Move keyboard focus between dialogs (non-modal) and between management windows.

Table 27. DS4000 Storage Manager alternate keyboard operations (continued)

Short cut	Action
Alt+ underlined letter	<p>Access menu items, buttons, and other interface components by using the keys associated with the underlined letters.</p> <p>For the menu options, select the Alt + underlined letter combination to access a main menu, and then select the underlined letter to access the individual menu item.</p> <p>For other interface components, use the Alt + underlined letter combination.</p>
Ctrl+F1	Display or conceal a tool tip when keyboard focus is on the toolbar.
Spacebar	Select an item or activate a hyperlink.
Ctrl+Spacebar (Contiguous/Non-contiguous) AMW Logical/Physical View	<p>Select multiple drives in the Physical View.</p> <p>To select multiple drives, select one drive by pressing Spacebar, and then press Tab to switch focus to the next drive you want to select; press Ctrl+Spacebar to select the drive.</p> <p>If you press Spacebar alone when multiple drives are selected then all selections are removed.</p> <p>Use the Ctrl+Spacebar combination to deselect a drive when multiple drives are selected.</p> <p>This behavior is the same for contiguous and non-contiguous selection of drives.</p>
End, Page Down	Move keyboard focus to the last item in the list.
Esc	Close the current dialog (does not require keyboard focus).
Home, Page Up	Move keyboard focus to the first item in the list.
Shift+Tab	Move keyboard focus through components in the reverse direction.
Ctrl+Tab	Move keyboard focus from a table to the next user interface component.
Tab	Navigate keyboard focus between components or select a hyperlink.
Down arrow	Move keyboard focus down one item in the list.
Left arrow	Move keyboard focus to the left.
Right arrow	Move keyboard focus to the right.
Up arrow	Move keyboard focus up one item in the list.

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 can 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

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 and abbreviations used in IBM TotalStorage DS4000 publications.

If you do not find the term you are looking for, see the *IBM Glossary of Computing Terms* located at the following Web site:

www.ibm.com/ibm/terminology

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.

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. One of three existing fibre-channel topologies, in which 2 - 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 a Stealth Mode.

arbitrated loop physical address (AL_PA). An 8-bit value that is used to uniquely identify an individual port within a loop. A loop can have one or more AL_PAs.

array. A collection of fibre-channel or SATA hard drives that are logically grouped together. All the drives in the array are assigned the same RAID level. An array is sometimes referred to as a "RAID set." See also *redundant array of independent disks (RAID)*, *RAID level*.

asynchronous write mode. In remote mirroring, an option that allows the primary controller to return a write I/O request completion to the host server before data has been successfully written by the secondary controller. See also *synchronous write mode*, *remote mirroring*, *Global Copy*, *Global Mirroring*.

AT. See *advanced technology (AT) bus architecture*.

ATA. See *AT-attached*.

AT-attached. Peripheral devices that are compatible with the original IBM AT computer standard in which signals on a 40-pin AT-attached (ATA) ribbon cable followed the timings and constraints of the Industry Standard Architecture (ISA) system bus on the IBM PC AT computer. Equivalent to integrated drive electronics (IDE).

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). The personal computer 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). In Transmission Control Protocol/Internet Protocol (TCP/IP) networking, an alternative protocol by which a diskless machine can obtain its Internet Protocol (IP) address and such configuration information as IP addresses of various servers from a BOOTP server.

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.

broadcast. The simultaneous transmission of data to more than one destination.

cathode ray tube (CRT). A display device in which controlled electron beams are used to display alphanumeric or graphical data on an electroluminescent screen.

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.

command. A statement used to initiate an action or start a service. A command consists of the command name abbreviation, and its parameters and flags if applicable. A command can be issued by typing it on a command line or selecting it from a menu.

community string. The name of a community contained in each Simple Network Management Protocol (SNMP) message.

concurrent download. A method of downloading and installing firmware that does not require the user to stop I/O to the controllers during the process.

CRC. See *cyclic redundancy check*.

CRT. See *cathode ray tube*.

CRU. See *customer replaceable unit*.

customer replaceable unit (CRU). An assembly or part that a customer can replace in its entirety when any of its components fail. Contrast with *field replaceable unit (FRU)*.

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*.

data striping. See *striping*.

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, or storage.

DHCP. See *Dynamic Host Configuration Protocol*.

direct access storage device (DASD). A device in which access time is effectively independent of the location of the data. Information is entered and retrieved without reference to previously accessed data. (For example, a disk drive is a DASD, in contrast with a tape drive, which stores data as a linear sequence.) DASDs include both fixed and removable storage devices.

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 Internet Protocol (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.

ECC. See *error correction coding*.

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

EISA. See *Extended Industry Standard Architecture*.

electrically erasable programmable read-only memory (EEPROM). A type of memory chip which can retain its contents without consistent electrical power. Unlike the PROM which can be programmed only once, the EEPROM can be erased electrically. Because it can only be reprogrammed a limited number of times before it wears out, 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.

environmental service module (ESM) canister. A component in a storage expansion enclosure that monitors the environmental condition of the components in that enclosure. Not all storage subsystems have ESM canisters.

E_port. See *expansion port*.

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

ESD. See *electrostatic discharge*.

ESM canister. See *environmental service module canister*.

EXP. See *storage expansion enclosure*.

expansion port (E_port). A port that connects the switches for two fabrics.

Extended Industry Standard Architecture (EISA). A bus standard for IBM compatibles that extends the Industry Standard Architecture (ISA) bus architecture to 32 bits and allows more than one central processing unit (CPU) to share the bus. See also *Industry Standard Architecture*.

fabric. A Fibre Channel entity which interconnects and facilitates logins of N_ports attached to it. The fabric is responsible for routing frames between source and destination N_ports using address information in the

frame header. A fabric can be as simple as a point-to-point channel between two N_ports, or as complex as a frame-routing switch that provides multiple and redundant internal pathways within the fabric between F_ports.

fabric port (F_port). In a fabric, an access point for connecting a user's N_port. An F_port facilitates N_port logins to the fabric from nodes connected to the fabric. An F_port is addressable by the N_port connected to it. See also *fabric*.

FAST MSJ. See *FAST Management Suite Java*.

FAST Management Suite Java (FAST 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.

FC. See *fibre channel*.

FC-AL. See *arbitrated loop*.

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 set of standards for a serial input/output (I/O) bus capable of transferring data between two ports at up to 100 Mbps, with standards proposals to go to higher speeds. FC supports point-to-point, arbitrated loop, and switched topologies.

Fibre Channel-Arbitrated Loop (FC-AL). See *arbitrated loop*.

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 commands, 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 field replaceable unit might contain other field replaceable units. Contrast with *customer replaceable unit (CRU)*.

FlashCopy. A premium feature for DS4000 that can make an instantaneous copy of the data in a volume.

F_port. See *fabric port*.

FRU. See *field replaceable unit*.

GBIC. See *gigabit interface converter*

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*.

Global Copy. Refers to a remote logical drive mirror pair that is set up using asynchronous write mode without the write consistency group option. This is also referred to as "Asynchronous Mirroring without Consistency Group." Global Copy does not ensure that write requests to multiple primary logical drives are carried out in the same order on the secondary logical drives as they are on the primary logical drives. If it is critical that writes to the primary logical drives are carried out in the same order in the appropriate secondary logical drives, Global Mirroring should be used instead of Global Copy. See also *asynchronous write mode, Global Mirroring, remote mirroring, Metro Mirroring*.

Global Mirroring. Refers to a remote logical drive mirror pair that is set up using asynchronous write mode with the write consistency group option. This is also referred to as "Asynchronous Mirroring with Consistency Group." Global Mirroring ensures that write requests to multiple primary logical drives are carried out in the same order on the secondary logical drives as they are on the primary logical drives, preventing data on the secondary logical drives from becoming inconsistent with the data on the primary logical drives. See also *asynchronous write mode, Global Copy, remote mirroring, Metro Mirroring*.

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 representing 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 (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.

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

host computer. See *host*.

host group. 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. Ports that physically reside on the host adapters and are automatically discovered by the DS4000 Storage Manager software. To give a host computer access to a partition, its associated host ports must be defined.

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 device driver that is used in a Novell NetWare environment to provide multipath input/output (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). Unofficial name for the bus architecture of the IBM PC/XT personal computer. This bus design included expansion slots for plugging in various adapter boards. Early versions had an 8-bit data path, later expanded to 16 bits. The "Extended Industry Standard Architecture" (EISA) further expanded the data path to 32 bits. See also *Extended Industry Standard Architecture*.

initial program load (IPL). The initialization procedure that causes an operating system to commence operation. Also referred to as a system restart, system startup, and boot.

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). 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. Also known as an Advanced Technology Attachment Interface (ATA).

Internet Protocol (IP). A protocol that routes data through a network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical network.

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.

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

IP. See *Internet Protocol*.

IPL. See *initial program load*.

IRQ. See *interrupt request*.

ISA. See *Industry Standard Architecture*.

Java Runtime Environment (JRE). A subset of the Java Development Kit (JDK) for end users and developers who want to redistribute the Java Runtime Environment (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 address*.

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

logical block address (LBA). The address of a logical block. Logical block addresses are typically used in hosts' I/O commands. The SCSI disk command protocol, for example, uses logical block addresses.

logical partition (LPAR). (1) A subset of a single system that contains resources (processors, memory, and input/output devices). A logical partition operates as an independent system. If hardware requirements are met, multiple logical partitions can exist within a system. (2) A fixed-size portion of a logical volume. A logical partition is the same size as the physical partitions in its volume group. Unless the logical volume of which it is a part is mirrored, each logical partition corresponds to, and its contents are stored on, a single physical partition. (3) One to three physical partitions (copies). The number of logical partitions within a logical volume is variable.

logical unit number (LUN). An identifier used on a small computer system 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 port. A node port (N_port) or fabric port (F_port) that supports arbitrated loop functions associated with an arbitrated loop topology.

LPAR. See *logical partition*.

LUN. See *logical unit number*.

MAC. See *medium access control*.

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

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.

MCA. See *micro channel architecture*.

media scan. A media scan is a background process that runs on all logical drives in the storage subsystem for which it has been enabled, providing error detection on the drive media. The media scan process scans all logical drive data to verify that it can be accessed, and optionally scans the logical drive redundancy information.

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.

Metro Mirroring. This term is used to refer to a remote logical drive mirror pair which is set up with synchronous write mode. See also *remote mirroring*, *Global Mirroring*.

MIB. See *management information base*.

micro channel architecture (MCA). Hardware that is used for PS/2 Model 50 computers and above to provide better growth potential and performance characteristics when compared with the original personal computer design.

Microsoft Cluster Server (MSCS). MSCS, a feature of Windows NT Server (Enterprise Edition), supports the connection of two servers into a cluster for higher availability and easier manageability. MSCS can automatically detect and recover from server or

application failures. It can also be used to balance server workload and provide for planned maintenance.

mini hub. An interface card or port device that receives short-wave fiber channel GBICs or SFPs. These devices enable redundant fibre channel connections from the host computers, either directly or through a fibre channel switch or managed hub, over optical fiber cables to the DS4000 Storage Server controllers. Each DS4000 controller is responsible for two mini hubs. Each mini hub has two ports. Four host ports (two on each controller) provide a cluster solution without use of a switch. Two host-side mini hubs are shipped as standard. See also *host port*, *gigabit interface converter (GBIC)*, *small form-factor pluggable (SFP)*.

mirroring. A fault-tolerance technique in which information on a hard disk is duplicated on additional hard disks. See also *remote mirroring*.

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

MSCS. See *Microsoft Cluster Server*.

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

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

N_port. See *node port*.

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.

partitioning. See *storage partition*.

parity check. (1) A test to determine whether the number of ones (or zeros) in an array of binary digits is odd or even. (2) A mathematical operation on the numerical representation of the information communicated between two pieces. For example, if parity is odd, any character represented by an even number has a bit added to it, making it odd, and an information receiver checks that each unit of information has an odd value.

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

PDF. See *portable document format*.

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

peripheral component interconnect local bus (PCI local bus). A local bus for PCs, from Intel, that provides a high-speed data path between the CPU and up to 10 peripherals (video, disk, network, and so on). The PCI bus coexists in the PC with the Industry Standard Architecture (ISA) or Extended Industry Standard Architecture (EISA) bus. ISA and EISA boards plug into an IA or EISA slot, while high-speed PCI controllers plug into a PCI slot. See also *Industry Standard Architecture*, *Extended Industry Standard Architecture*.

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

port. A part of the system unit or remote controller to which cables for external devices (such as display stations, terminals, printers, switches, or external storage units) are attached. The port is an access point for data entry or exit. 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. See also *arbitrated loop*.

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*.

RAID. See *redundant array of independent disks (RAID)*.

RAID level. An array's RAID level is a number that refers to the method used to achieve redundancy and fault tolerance in the array. See also *array, redundant array of independent disks (RAID)*.

RAID set. See *array*.

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. Contrast with *DASD*.

RDAC. See *redundant disk array controller*.

read-only memory (ROM). Memory in which stored data cannot be changed by the user 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.

redundant array of independent disks (RAID). A collection of disk drives (*array*) that appears as a single volume to the server, which is fault tolerant through an assigned method of data striping, mirroring, or parity checking. Each array is assigned a RAID level, which is a specific number that refers to the method used to achieve redundancy and fault tolerance. See also *array, parity check, mirroring, RAID level, striping*.

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.

remote mirroring. Online, real-time replication of data between storage subsystems that are maintained on separate media. The Enhanced Remote Mirror Option is a DS4000 premium feature that provides support for remote mirroring. See also *Global Mirroring, Metro Mirroring*.

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*.

SAI. See *Storage Array Identifier*.

SA Identifier. See *Storage Array Identifier*.

SAN. See *storage area network*.

SATA. See *serial ATA*.

scope. Defines a group of controllers by their Internet Protocol (IP) addresses. A scope must be created and defined so that dynamic IP addresses can be assigned to controllers on the network.

SCSI. See *small computer system interface*.

segmented loop port (SL_port). A port that allows division of 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 ATA. The standard for a high-speed alternative to small computer system interface (SCSI) hard drives. The SATA-1 standard is equivalent in performance to a 10 000 RPM SCSI drive.

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 DS4000 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 DS4000 Storage Manager client software, which is a Java-based graphical user interface (GUI) that is used to configure, manage, and troubleshoot storage servers and storage expansion enclosures in a DS4000 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 DS4000 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 cached data of the operating system 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* and *SNMPv1*.

SNMP trap event. (1) (2) An event notification sent by the SNMP agent that identifies conditions, such as thresholds, that exceed a predetermined value. See also *Simple Network Management Protocol*.

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

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 turned off.

storage area network (SAN). A dedicated storage network tailored to a specific environment, combining

servers, storage products, networking products, software, and services. See also *fabric*.

Storage Array Identifier (SAI or SA Identifier). The Storage Array Identifier is the identification value used by the DS4000 Storage Manager host software (SMClient) to uniquely identify each managed storage server. The DS4000 Storage Manager SMClient program maintains Storage Array Identifier records of previously-discovered storage servers in the host resident file, which allows it to retain discovery information in a persistent fashion.

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

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 (I/O) path.

storage partition. Storage subsystem logical drives that are visible to a host computer or are shared among host computers that are part of a host group.

storage partition topology. In the DS4000 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. The host port, host computer, and host group topological elements must be defined to grant access to host computers and host groups using logical drive-to-LUN mappings.

striping. Splitting data to be written into equal blocks and writing blocks simultaneously to separate disk drives. Striping maximizes performance to the disks. Reading the data back is also scheduled in parallel, with a block being read concurrently from each disk then reassembled at the host.

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

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 in the network.

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 zoning. See *zoning*.

synchronous write mode. In remote mirroring, an option that requires the primary controller to wait for the

acknowledgment of a write operation from the secondary controller before returning a write I/O request completion to the host. See also *asynchronous write mode*, *remote mirroring*, *Metro Mirroring*.

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.

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 to transmit and receive data. Transceiver is an abbreviation of transmitter-receiver.

translated loop port (TL_port). A port that 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 packet-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.

worldwide port name (WWPN). A unique identifier for a switch on local and global networks.

worldwide name (WWN). A globally unique 64-bit identifier assigned to each Fibre Channel port.

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

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.

WWN. See *worldwide name*.

zoning. (1) In Fibre Channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones. (2) A function that allows segmentation of nodes by address, name, or physical port and is provided by fabric switches or hubs.

Index

A

- about this document xi
- access volume 24
- add Storage Subsystem option 28
- address, IBM xvii
- ADT feature 48, 50
- Advanced menu 35
- AIX and Sun Solaris, failover protection 50
- alert destinations 75
 - configuration 85
 - configuring 85
 - setting 85
- alert notification
 - configuring alert destinations 85
 - mail server configuration 85
 - overview 84
 - selecting the node 85
 - setting 86
 - setting alert destinations 85
- array 12, 47, 51
- Array menu 33
- asynchronous write mode 63
- audience xi
- Auto-Logical Drive Transfer (ADT) feature 48
- automatic discovery option 28

B

- background media scan 55

C

- cache flush
 - described 54
 - performance impacts 54
 - settings 54
 - start percentage 54
 - stop flush percentage 54
- cache hit
 - optimizing 95
 - percentage 95
- cache read-ahead, choosing a multiplier 94
- capacity
 - Dynamic Capacity Expansion (DCE) 48
 - free 12
 - free and unconfigured 66
 - unconfigured 12
- channel protection, using 69
- coexisting storage subsystems, managing 26
- command line interface (SMcli)
 - examples 43
 - overview 38
 - parameters 39
 - usage and formatting requirements 42
 - using 38
- comments, how to send xvii

- components
 - button 91
 - software 14
 - storage subsystem 10
- configuration
 - mail server 85
 - sender address 85
 - storage subsystem 69
- Contacting Device status 79
- controller
 - cache memory, data protection 53
 - description 11
 - enclosure 11, 54
 - transfer rate, optimizing 94
- Controller menu 34
- copy services
 - Enhanced Remote Mirroring option 59
 - FlashCopy 59
 - VolumeCopy 59
- copyback 54
- critical event
 - notification 85
 - problem solving 97
- customer support alert notification
 - how to configure 85

D

- data
 - backing up 62
 - copying for greater access 62
 - path failover protection 48
 - protection 116
 - protection in the controller cache memory 53
 - protection strategies 45
 - redundancy 51
 - restoring FlashCopy logical drive data 62
- DCE (Dynamic Capacity Expansion) 48
- default host group, defined 71
- default logical drive-to-LUN mapping
 - defined 72
- default LUN 71
- default settings for failover protection 50
- device drivers
 - downloading latest versions 1
- Device Table 78
- Device Tree 78
- DHCP/BOOTP server 22, 25
- direct (out-of-band) management method
 - advantages 21
 - described 21
 - disadvantages 22
- directly managed storage subsystems 21
- disk access, minimize 96
- document organization xiv
- documentation
 - DS4000 119
 - DS4000-related documents 124

- documentation (*continued*)
 - DS4100 SATA Storage Server 122
 - DS4300 Fibre Channel Storage Server 121
 - DS4400 Fibre Channel Storage Server 120
 - DS4500 Fibre Channel Storage Server 119
 - DS4800 Storage Subsystem 123
 - using xv
- drive 11
- drive firmware, downloading 81
- Drive menu 34
- drivers
 - See device drivers
- drives, logical 12, 45
- DS4000 documentation 119
- DS4000 Hardware Maintenance Manual 124
- DS4000 Problem Determination Guide 124
- DS4000 Quick Start Guide 124
- DS4000 Storage Manager
 - related documents 124
- DS4000/FAST product renaming 2
- DS4100 SATA Storage Server library 122
- DS4300 Fibre Channel Storage Server library 121
- DS4400 Storage Server library 120
- DS4500 Fibre Channel Storage Server library 119
- DS4800 Storage Subsystem library 123
- DVE (Dynamic Logical Drive Expansion) 47
- Dynamic Capacity Expansion (DCE) 48
- Dynamic Logical Drive Expansion (DVE) 47

E

- emwdata.bin file 27
- Enhanced Remote Mirroring option
 - asynchronous write mode 63
 - description 59
 - diagnostics 64
 - enhancements 63
 - introduction 63
 - logical drive types 65
 - mirror relationships 66
 - mirror repository logical drives 66
 - number of mirror relationships per subsystem 64
 - primary logical drives 65
 - read access 64
 - resynchronization methods 64
 - role reversal 65
 - secondary logical drives 65
 - suspend and resume 63
 - write modes 66
 - write order consistency 63
- enhancements 6
- Enterprise Management window 27
 - component of SMclient 14
 - Device Table 78
 - Device Tree 28, 78
 - Help 1
 - maintaining storage subsystems 78
 - monitoring storage subsystems 75, 78
 - Needs Attention icon 79
 - overall health status pane 78
 - status icons displayed in 78

- Enterprise Management window (*continued*)
 - synchronizing 87
- Environmental Services Module (ESM) card 83
- errors, media scan 57
- ESM
 - downloading card firmware 83
 - overview 83
- event log 97, 116
- Event Monitor
 - and Enterprise Management window, synchronizing 87
 - example 87
 - installing 86
 - overview 86
 - setting alert notifications 86
 - synchronizing the Enterprise Management window 87
- event notification 85, 116
- examples, SMcli 43

F

- fabric switches 19
- failover protection
 - ADT feature 48
 - AIX and Sun Solaris 50
 - default settings 50
 - HP-UX 50
 - Linux 50
 - Microsoft Windows 49
 - Novell NetWare 49
 - operating system specific 49
 - overview 48
 - RDAC feature 49
- failure notification
 - example 79
 - in the Subsystem Management window 79
- failures, recovering from 89
- FAST-MSJ
 - failover driver 18
- FAST/DS4000 product renaming 2
- feature key, obtaining 72
- features, new 6
- Fibre Channel I/O
 - access pattern 94
 - balancing the load 93
 - request rate optimizing 94
 - size 94
- Fibre Channel switches 20
- files, defragmenting 96
- firmware
 - downloading 80
 - new features
 - version 6.10.xx.xx 8
 - version 6.12.xx.xx 7
 - updating in the storage expansion enclosures 79
 - updating in the storage subsystem 79
 - version 5
- Fixing status 78
- FlashCopy
 - description 59

- FlashCopy (*continued*)
 - logical drive 45
 - overview 10, 61
 - repository logical drive 45
 - script scenarios 59
- form, reader comment xvii
- free capacity 12, 66
- free-capacity nodes 69
- full synchronization 65

G

- Global Copy 63
- Global Mirroring 63
- glossary 131
- graphical user interface (GUI)
 - managing the storage subsystem 27

H

- hardware
 - requirements 19
- hardware components
 - DHCP server, BOOTP or BOOTP compliant 20
 - file server 21
 - host computer 20
 - management station 20
 - network-management station 20
 - storage subsystem 20
- hardware service and support xvii
- Help menu 35
- help, getting xv
- heterogeneous hosts
 - defining types 72
 - overview 72
- host adapters 19
- host bus adapters 20
- host computer 6, 71
- host group
 - definition 70
 - description 71
- host port
 - defined 71, 72
 - discovery of 71
- host-agent managed storage subsystems 23
- host-agent management method
 - advantages 23
 - described 23
 - disadvantages 24
- Hot Add utility 18
- hot spare drive
 - configuring 54
 - defined 54
- how to send your comments xvii
- HP-UX, failover protection 50

I

- I/O access pattern and I/O size 94
- I/O data field 93
- I/O data path protection 48

- I/O request rate
 - impact from cache flush settings 54
 - optimizing 94
- I/O transfer rate, optimizing 94
- IBM address xvii
- IBM Safety Information 124
- IBM Storage Area Network (IBMSAN) driver 17
- information from the World Wide Web xvi
- information, getting xv
- Intermix
 - enabling with NVSRAM (firmware 6.10.xx.xx) 9
 - enabling with premium feature key 7

L

- Linux
 - failover driver 18
 - failover protection 50
- local storage subsystems 63
- Logical Drive menu 33
- logical drive types
 - primary 65
 - secondary 65
- logical drive-to-LUN mapping
 - default 72
 - defined 71
 - specific 71
- logical drive-to-LUN terminology
 - default host group 71
 - host 71
 - host group 71
 - host port 71
 - mapping 71, 72
 - storage partition topology 70
 - storage partitions 70
 - mapping preference 72
- logical drives
 - base 61
 - creating step-by-step 69
 - definition 12
 - Dynamic Logical Drive Expansion (DVE) 47
 - FlashCopy 45, 61
 - FlashCopy repository 45
 - mirror relationship 66
 - mirror repository 46, 66
 - missing 83
 - modification priority setting 95
 - overview 45
 - primary 46
 - recovering 83
 - repository 61
 - secondary 46
 - source 46
 - standard 45
 - target 46
 - VolumeCopy 62
- Logical/Physical view 29, 78
- Logical/Physical View 31
- LUN
 - address space 71
 - defined 71

M

- machine types and supported software 3
- mail server configuration 85
- managed hub 19
- management domain, populating 113
 - automatic discovery option 28
 - overview 28
 - using Add Storage Subsystem 28
- management methods for storage subsystem
 - direct (out-of-band) management method 21
 - host-agent management method 23
- management station 6, 19
- management, storage subsystem
 - direct (out-of-band) 21
 - host-agent 23
 - overview 21
- Mappings menu 33
- Mappings View 31
- media scan
 - changing settings 55
 - duration 58
 - errors reported 57
 - overview 55
 - performance impact 56
 - settings 58
- medical imaging applications 53
- menus, Subsystem Management window 31
- Microsoft Windows failover protection 49
- mirror relationships 66
- mirror repository 65
- mirror repository logical drives 46, 66
- missing logical drives, viewing and recovering 83
- multi-user environments 53
- multimedia applications 53

N

- Needs Attention
 - icon 79
 - status 78
- new features 6
- new features in this edition 6
- notes, important 130
- notices xv, 129
- notification
 - alert 84
 - configuring alert destinations 85
 - failure 79
 - of events 116
 - selecting the node 85
 - setting alert destinations 85
 - setting alert notifications 86
- Novell NetWare failover protection 49
- NVSRAM, downloading
 - from a firmware image 81
 - from a standalone image 81

O

- online help systems
 - configuring storage partitions 115
 - configuring storage subsystems 113, 114
 - Enterprise Management window 115
 - event notification 116
 - miscellaneous system administration 117
 - performance and tuning 118
 - populating a management domain 113
 - protecting data 116
 - recovering from problems 117
 - security 118
 - Subsystem Management window 113, 114, 115
 - using a script editor 114
- operating system specific failover protection 49
- organization of the document xiv
- overall health status 78
- ownership, preferred controller 51

P

- parallel drive firmware download 82
- parameters, SMcli 39
- parity 51
- password protection, configuring 67
- performance and tuning 118
- performance monitor 93
- Persistent Reservations, managing 67
- physical view, subsystem-management window 29
- point-in-time (PIT) image 61
- power outage 54
- preferred controller ownership 51
- premium feature support
 - restrictions 60
- premium features
 - Enhanced Remote Mirroring option 59
 - FlashCopy 59
 - Intermix 7, 9
 - VolumeCopy 59
- primary logical drive 46
- primary logical drives 65
- priority setting, modification 95
- problem recovery 117
- problem solving, critical event 97

Q

- QLremote 18
- quick reference status
 - Contacting Device 79
 - Fixing 78
 - Optimal 78
 - Optimal status 78
 - Unresponsive 79

R

- RAID level
 - and channel protection 69
 - application behavior 53, 95

- RAID level (*continued*)
 - choosing 53, 95
 - configurations 52
 - data redundancy 51
 - described 51
- RAID-0
 - described 52
 - drive failure consequences 52
- RAID-1
 - described 52
 - drive failure consequences 52
- RAID-3
 - described 53
 - drive failure consequences 53
- RAID-5
 - described 53
 - drive failure consequences 53
- RDAC feature 49
- reader comment form processing xvii
- reconstruction 54
- Recovery Guru
 - Recovery Procedure 89
 - Summary area 89
 - window 88
- redundancy of Fibre Channel arbitrated loops 69
- Redundant disk array controller (RDAC) 16
- reference, task 113
- remote mirror setup, logical drive types 65
- remote storage subsystems 63
- renaming 2
- requirements
 - hardware 19
 - SMcli 42
- restrictions
 - premium feature support 60
- resynchronization methods 64
- role reversal 65

S

- sample network, reviewing 25
- script editor
 - adding comments to a script 37
 - using 36, 114
 - window 35
- secondary logical drive 46, 65
- security 118
- segment size, choosing 96
- sender address configuration 85
- service and support
 - hardware xvii
 - software xvii
- settings, media scan 58
- Simple Network Management Protocol (SNMP)
 - traps 25
- SMagent disk space requirements 15
- SMcli
 - examples 43
 - overview 38
 - parameters 39
 - usage and formatting requirements 42

- SMcli (*continued*)
 - using 38
- SMclient 14
- SMdevices utility 18
- software components
 - RDAC 16
 - SMagent 15
 - SMclient 14
- software service and support xvii
- software, supported 3
- source logical drive 46
- staged controller firmware download 81
- starpubs xvii
- start percentage, cache flush 54
- stop percentage, cache flush 54
- storage area network (SAN)
 - technical support Web site xvi
- storage expansion enclosure 11
- storage expansion enclosures, updating the firmware 79
- storage management software
 - Enterprise Management window 27
 - hardware requirements
 - BOOTP server 19
 - installation requirements 19
 - new terminology 5
 - Subsystem Management window 29
- Storage Manager 9.1 client 14
- Storage Manager software
 - new features 7, 8
- Storage Manager Utility (SMutil) 18
- storage partition
 - feature 10, 72
 - switch zoning 70
- storage partition topology, defined 70
- storage partitioning specifications 12
- storage partitions
 - configuring 115
 - creating 70
 - described 70
 - description 12
 - enabled 12
 - feature key 72
 - major steps to creating 86
- storage subsystem
 - components 10
 - configuration 69, 114
 - creating logical drives 69
 - description 19
 - device tree 28
 - failure notification 79
 - hardware requirements 19
 - logical components 12
 - maintaining and monitoring 75
 - maintaining in a management domain 78
 - managing using the graphical user interface 27
 - password protection configuration 67
 - physical components 11
 - quick reference status icon 78
 - status quick reference 78
 - tuning options available 93

- storage subsystem (*continued*)
 - updating the firmware 79
- storage subsystem management
 - direct (out-of-band) 21
 - host-agent 23
 - overview 21
- Storage Subsystem menu 32
- storage subsystems
 - coexisting 26
 - directly managed 21
 - host-agent managed 23
 - local and remote 63
 - tuning 93
- storage subsystems maintenance
 - in a management domain 78
 - overview 75
- storage-partition mapping preference
 - defined 72
- storage-subsystem failures, recovering from 89
- Subsystem Management window 79
 - Advanced menu 35
 - Array menu 33
 - component of SMclient 14
 - Controller menu 34
 - Drive menu 34
 - event log 97
 - Help 1
 - Help menu 35
 - Logical Drive menu 33
 - Logical/Physical View 31
 - Mappings menu 33
 - Mappings View 31
 - menus 31
 - monitoring storage subsystems with 75
 - overview 29
 - Storage Subsystem menu 32
 - tabs 30
 - View menu 32
- supported software 3
- suspend and resume mirror synchronization 63
- switch
 - technical support Web site xvi
 - zoning 70
- system administration 117

T

- target logical drive 46
- task reference 113
- tasks by documentation title 119
- technical assistance xv
- terminology 5
- topological elements, when to define 70
- trademarks 129
- transfer rate 93

U

- unconfigured capacity 12, 66
- unconfigured nodes 69
- UNIX BOOTP server 19

- Unresponsive status 79

V

- version, firmware 5
- View menu 32
- VolumeCopy
 - backing up data 62
 - copying data for greater access 62
 - description 59
 - overview 10, 62
 - restoring FlashCopy logical drive data to the base logical drive 62

W

- Web sites
 - DS4000 interoperability matrix xvi
 - DS4000 storage subsystems xvi
 - DS4000 technical support xvi
 - IBM publications center xvii
 - IBM TotalStorage products xvi
 - premium feature activation xvi
 - readme files xvi
 - SAN support xvi
 - switch support xvi
- who should read this document xi
- window, script editor 35
- World Wide Web, getting help and information xvi
- write cache mirroring
 - described 53
 - how to enable 53
- write caching
 - and data loss 54
 - and performance 53
 - enabling 95
- write order consistency 63

Z

- zoning 70

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

IBM TotalStorage DS4000 Storage Manager Version 9
Concepts Guide

Publication No. GC26-7734-02

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department GZW
9000 South Rita Road
Tucson, Arizona
U.S.A 85744-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GC26-7734-02

