

IBM TotalStorage FAStT Storage Manager Version 8.3



Installation and Support Guide for AIX, HP-UX, and Solaris

IBM TotalStorage FAStT Storage Manager Version 8.3



Installation and Support Guide for AIX, HP-UX, and Solaris

Note:

Before using this information and the product it supports, read the information in "Notices" on page 151.

Second Edition (April 2003)

This edition replaces GC26-7521-00

© **Copyright International Business Machines Corporation 2003. 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
FAST installation process overview	xi
FAST documentation	xii
FAST Storage Manager Version 8.3 library	xiii
FAST900 Fibre Channel Storage Server library.	xiv
FAST700 Fibre Channel Storage Server library	xv
FAST600 Fibre Channel Storage Server library.	xvi
FAST500 Fibre Channel Storage Server library	xvii
FAST200 Fibre Channel Storage Server library	xviii
FAST related documents	xix
How this document is organized	xx
Online help	xx
Web sites.	xxi
How to send your comments.	xxi
Chapter 1. Introduction	1
New system features	2
Storage Manager 8 premium features	2
FAST600 restrictions	3
Storage Manager 8 software packages	3
Storage Manager 8 client software package.	4
Storage Manager 8 agent software package	4
Storage Manager 8 utility software package.	5
RDAC	5
Storage subsystem management methods	5
Host-agent (in-band) management method	5
Direct (out-of-band) management method	7
Operating system requirements	8
Setting up IP addresses for FAST storage controllers	8
Enabling Storage Partitioning.	10
Obtaining the feature enable identifier	10
Generating the feature key file and enabling Storage Partitioning	10
Chapter 2. Installing storage management station software on AIX systems 13	
Hardware and firmware requirements	13
AIX restrictions	14
Installing the client software	15
Prerequisites.	15
Installing SMruntime	16
Installing SMclient.	16
Performing the initial configuration of storage subsystems on AIX hosts	16
Updating FAST firmware and NVSRAM	17
Setting up an AIX host group.	18
Creating direct and SAN-attached configurations	19
Creating a direct-attached configuration	19
Creating a SAN-attached configuration	19
Installing the RDAC driver.	20
Prerequisites.	20

Installing the RDAC driver	20
AIX configuration information	21
Viewing and setting attributes of the RDAC driver for AIX	21
Viewing Object Data Manager (ODM) attributes in AIX	26
Changing ODM attribute settings in AIX	27
Verifying the installation and configuration of AIX hosts	28
Identifying the controller ID numbers	29
Identifying device names and bus numbers	30
Identifying logical drives by operating system device names	33
Identifying FlashCopy volume types	35
Using dynamic capacity expansion and dynamic volume expansion with AIX	36
Performing a dynamic capacity expansion operation	36
Performing a dynamic volume expansion (DVE) operation	36
Resolving disk array errors	37
Redistributing volumes in case of failure	39

Chapter 3. Installing storage management station software on HP-UX systems

Hardware and firmware requirements	41
Installing the client software	42
Prerequisites.	42
Installing SMruntime	43
Installing the SMclient	43
Installing the host software	44
Prerequisites.	44
Installing the SMagent (optional)	44
Installing SMutil.	45
Performing the initial configuration of storage subsystems on HP-UX hosts	46
Steps for adding storage subsystems to SMclient	46
Steps for updating NVSRAM and FAStT firmware	47
Setting up an HP-UX host group	47
Creating direct-attached and SAN-attached configurations	48
Creating a direct-attached configuration	48
Creating a SAN-attached configuration	48
Enabling multipath I/O with PV-links	49
Identifying the controller ID numbers	51
Identifying device names and bus numbers	52
Redistributing volumes in case of failure	53

Chapter 4. Installing storage management station software on Solaris systems

Hardware and firmware requirements	55
Installing the client software	57
Prerequisites.	57
Installing SMruntime	57
Installing SMclient.	58
Installing host software on Solaris hosts.	59
Prerequisites.	59
Installing SMagent (optional)	59
Installing SMutil.	60
Uninstalling Storage Manager 8.3 software	61
Performing the initial configuration of storage subsystems on Solaris hosts	61
Steps for adding storage subsystems to SMclient	62
Steps for updating NVSRAM and FAStT firmware	62
Steps for setting up a Solaris host group	62
Default partitioning for Solaris devices	63

Direct-attached and SAN-attached configurations	64
Creating a direct-attached configuration	64
Creating a SAN-attached configuration	64
Configuring JNI host bus adapter cards	65
Steps for installing the JNI adapter driver package	65
Steps for modifying JNI HBA settings.	66
Installing RDAC on Solaris hosts	69
Verifying external storage	71
Redistributing volumes in case of failure	72
Chapter 5. Completing the software installation.	73
Failover protection	73
Changing the host type	73
Understanding heterogeneous hosts	74
Configuring storage subsystems	74
Starting the Enterprise Management window	74
Configuring alert notifications.	77
Starting the Subsystem Management window	77
Renaming storage subsystems	78
Performing other storage subsystem management tasks	78
Chapter 6. Storage Manager 8.3 with high-availability cluster services.	81
General information	81
Prerequisites for HP-UX	81
Prerequisites for Solaris and Veritas Cluster Server	81
General hardware requirements.	82
System dependencies	82
Storage Manager 8.3 with High Availability Cluster Management Program (HACMP)	83
FAST software and firmware requirements	83
HACMP software requirements	83
Storage Manager 8.3 with Parallel System Support Programs (PSSP) and General Parallel File System (GPFS)	84
PSSP, and GPFS software requirements	84
Configuration limitations	84
Other usage notes specific to PSSP/GPFS environments	85
GPFS/PSSP logical configuration diagrams	85
Chapter 7. Copy Services and the FAST Storage Server	93
Overview of FlashCopy	93
Enabling FlashCopy	94
Obtaining the feature enable identifier	94
Generating the feature key file and enabling FlashCopy	94
Planning a FlashCopy logical drive	95
Setting FlashCopy repository logical drive failure options	95
Viewing current FlashCopy repository logical drive failure settings	95
Estimating FlashCopy repository logical drive capacity	98
Estimating FlashCopy repository life	99
Creating a FlashCopy logical drive	100
Creating FlashCopy logical drives on AIX, HP-UX, and Solaris	110
Using the script editor and command-line interface	121
Checking FlashCopy status	121
Mapping a FlashCopy drive to a host	123
Viewing the FlashCopy logical drive status	126
Disabling a FlashCopy logical drive	131
Recreating a FlashCopy logical drive	133

Resizing a FlashCopy repository logical drive	134
Deleting a FlashCopy drive	136
Viewing and recovering missing logical drives	138
Appendix A. AIX system requirements	139
Hardware requirements	139
Software requirements	139
RDAC installation requirements	139
Appendix B. HP-UX system requirements	141
Appendix C. Solaris system requirements	143
Appendix D. MC/Service Guard configuration details	145
Appendix E. JNI host bus adapter settings	147
Notices	151
Trademarks.	151
Glossary	153
Index	161

Figures

1. Installation process flow by current publications.	xii
2. Host-agent (in-band) managed storage subsystems	6
3. Direct (out-of-band) managed storage subsystems	7
4. Controller Properties window	30
5. Use of commands applied to validating system configuration.	31
6. Change Mappings window	34
7. Volume group Blast50 is increased from 4 GB to 8 GB	37
8. Device names	50
9. Controller Properties window	52
10. Device identification information	53
11. Initial Automatic Discovery window	75
12. Enterprise Management window	76
13. Subsystem Management window	78
14. RVSD cluster configuration with single FASTT500/700/900 - Storage Manager 8.3, shown with two partitions (four maximum) per FASTT	86
15. RVSD cluster configuration with multiple FASTT500/700/900 - Storage Manager 8.3, shown with one partition per FASTT	87
16. Typical RVSD cluster configuration with multiple FASTT500/700/900 - Storage Manager 8.3, shown with one partition per FASTT.	88
17. RVSD cluster configuration with multiple FASTT500/700/900 - Storage Manager 8.3, shown with two partitions per FASTT	89
18. HACMP/GPFS cluster configuration with single FASTT500/700/900, shown with one partition per FASTT	90
19. HACMP/GPFS cluster configuration with multiple FASTT500/700/900, shown with two partitions per FASTT	91
20. FlashCopy repository logical drive properties	96
21. Repository capacity settings.	97
22. Create FlashCopy menu options.	101
23. Create a FlashCopy Logical Drive Wizard startup	102
24. Wizard Introduction window	103
25. Allocate Capacity window	104
26. Specify Logical Drive Parameters	105
27. Specify Array Parameters window	106
28. Specify Names window	107
29. Specify Repository Drive Capacity	108
30. Completed window.	109
31. FlashCopy volumes	110
32. Listing premium features	122
33. Features list showing FlashCopy enabled	122
34. Feature list showing FlashCopy not enabled	123
35. FlashCopy feature icon	123
36. FlashCopy feature icon disabled.	123
37. Undefined FlashCopy disk	124
38. Define Additional Mapping	124
39. Define Additional Mapping window	125
40. Mapped FlashCopy logical disk	126
41. FlashCopy icon states	127
42. Repository Logical Drive Properties	128
43. Base Repository Logical Drive Properties	129
44. Repository drive warning threshold.	130
45. Choosing to disable the FlashCopy drive	132
46. Disable FlashCopy confirmation window	132
47. Icon showing the disabled FlashCopy logical drive	133

48. Recreating a FlashCopy logical drive	134
49. Deleting the FlashCopy logical drive	137
50. Delete FlashCopy Logical Drive Dialog window	137

Tables

1.	TotalStorage FAStT Storage Manager Version 8.3 titles by user tasks	xiii
2.	TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks	xiv
3.	TotalStorage FAStT700 Fibre Channel Storage Server document titles by user tasks	xv
4.	TotalStorage FAStT600 Fibre Channel Storage Server document titles by user tasks	xvi
5.	TotalStorage FAStT500 and FAStT High Availablity Storage Server document titles by user tasks	xvii
6.	TotalStorage FAStT200 and FAStT High Availablity Storage Server document titles by user tasks	xviii
7.	TotalStorage FAStT related document titles by user tasks	xix
8.	Installation sequences of Storage Manager 8.3 software packages by host type	3
9.	FAStT storage servers that are supported on AIX systems	13
10.	Supported versions of hardware for AIX systems	13
11.	Attributes for dar devices	22
12.	Attributes for dac devices	23
13.	Attributes for hdisk devices	24
14.	Supported versions of hardware for HP-UX systems	41
15.	Primary and secondary path matrix information	51
16.	FAStT storage servers that are supported on Solaris systems	55
17.	Supported versions of hardware for Solaris systems	55
18.	JNI adapter configuration files	66
19.	HACMP software requirements	83
20.	Filesets required for AIX 4.3.3 RDAC	139
21.	Filesets required for AIX 5.1 RDAC	139
22.	Filesets required for AIX 5.2 RDAC	140
23.	HP-UX kernel parameter configuration requirements	142
24.	Configuration file name: /kernel/drv/fca-pci.conf	147
25.	Configuration file name: /kernel/drv/jnic146x.conf	147
26.	Configuration file name: /kernel/drv/jnic.conf	148
27.	Configuration file name: /kernel/drv/fcaw.conf	148

About this document

This document provides information about how to set up, install, configure, and work with the IBM TotalStorage™ Fibre Array Storage Technology (FAStT) Storage Manager Version 8.3 in AIX™, HP-UX and Solaris.

Throughout this document, the terms *storage management software* and *Storage Manager 8.3* refer to the IBM TotalStorage FAStT Storage Manager Version 8.3. Individual components of the storage management software are identified by name.

Note: Throughout this document, Storage Manager version 8.3 refers to software levels 8.30 or higher.

Use this document to:

- Determine the hardware and software that is required to install the management software into your subsystem network
- Install the management software
- Upgrade controller NVSRAM and firmware
- Identify management features that are unique to your specific installation

Note: This document does not cover hardware installation or integration. For information about these topics, see the FAStT Installation and Support Guide that is appropriate for your FAStT Fibre Channel Storage Server.


Who should read this document

This document is intended for system administrators and storage administrators who are responsible for installing software. Readers should have knowledge of RAID, SCSI, and fibre-channel technology, and should also have working knowledge of the applicable operating systems that are used with the management software.

FAStT installation process overview

The following flow chart gives an overview of the installation process for the FAStT hardware and the FAStT Storage Manager software. Lined arrows in the flow chart indicate consecutive steps in the hardware and software installation process. Labeled arrows indicate which current documents provide detailed information about

those steps.



SJ001046

Figure 1. Installation process flow by current publications

FASTt documentation

The following tables present an overview of the FASTt Storage Manager and the FASTt900, FASTt700, FASTt600, FASTt500, and FASTt200 Fibre Channel Storage Server document libraries, as well as related documents. Each table lists documents that are included in the libraries and where to locate the information that you need to accomplish common tasks.

FAST Storage Manager Version 8.3 library

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

Table 1. TotalStorage FAST Storage Manager Version 8.3 titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM TotalStorage FAST Storage Manager 8.3 Installation and Support Guide for Microsoft® Windows®NT and Windows 2000, GC26-7522	✓		✓	✓		
IBM TotalStorage FAST Storage Manager 8.3 Installation and Support Guide for Linux, GC26-7519	✓		✓	✓		
IBM TotalStorage FAST Storage Manager 8.3 Installation and Support Guide for Novell NetWare, GC26-7520	✓		✓	✓		
IBM TotalStorage FAST Storage Manager 8.3 Installation and Support Guide UNIX and AIX®Environments, GC26-7521	✓		✓	✓		
IBM FAST Remote Mirror Option Installation and User's Guide, 48P9821	✓		✓	✓	✓	
IBM FAST Storage Manager Script Commands (see product CD)				✓		
IBM FAST Storage Manager Version 7.10 Concepts Guide, 25P1661	✓	✓	✓	✓	✓	✓

FAST900 Fibre Channel Storage Server library

Table 2 associates each document in the FAST900 Fibre Channel Storage Server library with its related common user tasks.

Table 2. TotalStorage FAST900 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM TotalStorage FAST900 Installation and Support Guide, GC26-7530	✓	✓		✓		
IBM TotalStorage FAST900 Fibre Channel Cabling Instructions, 24P8135	✓	✓				
IBM TotalStorage FAST900 User's Guide, GC26-7534				✓	✓	✓
IBM TotalStorage FAST FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
IBM FAST FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
IBM TotalStorage FAST900 Rack Mounting Instructions, 19K0900	✓	✓				
IBM Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329	✓	✓			✓	✓
IBM FAST Management Suite Java User's Guide, 32P0081					✓	✓
IBM TotalStorage FAST Fibre Channel Hardware Maintenance Manual and Problem Determination Guide, GC26-7528						✓

FAStT700 Fibre Channel Storage Server library

Table 3 associates each document in the FAStT700 Fibre Channel Storage Server library with its related common user tasks.

Table 3. TotalStorage FAStT700 Fibre Channel Storage Server document titles by user tasks

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

FAStT600 Fibre Channel Storage Server library

Table 4 associates each document in the FAStT600 Fibre Channel Storage Server library with its related common user tasks.

Table 4. TotalStorage FAStT600 Fibre Channel Storage Server document titles by user tasks

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

FAStT500 Fibre Channel Storage Server library

Table 5 associates each document in the FAStT500 Fibre Channel Storage Server library with its related common user tasks.

Table 5. TotalStorage FAStT500 and FAStT High Availability Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM FAStT500 RAID Controller Enclosure Unit User's Guide, 48P9847				✓	✓	✓
IBM FAStT EXP500 Storage Expansion Unit Installation and User's Guide, 59P5637	✓	✓		✓	✓	✓
IBM FAStT500 RAID Controller Enclosure Unit Installation Guide, 59P6244		✓			✓	
IBM FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
IBM TotalStorage FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
IBM FAStT Management Suite Java User's Guide, 32P0081					✓	✓
IBM Fibre Channel Hardware Maintenance Manual, 19K6130						✓
IBM Fibre Channel Problem Determination Guide, 48P9804						✓

FAStT200 Fibre Channel Storage Server library

Table 6 associates each document in the FAStT200 Fibre Channel Storage Server library with its related common user tasks.

Table 6. TotalStorage FAStT200 and FAStT High Availability Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
IBM FAStT200 and FAStT200 HA Storage Servers Installation and User's Guide, 59P6243	✓	✓		✓	✓	
IBM FAStT200 Fibre Channel Cabling Instructions, 21P9094	✓	✓				
IBM FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
IBM FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
IBM FAStT Management Suite Java User's Guide, 32P0081					✓	✓
IBM Fibre Channel Hardware Maintenance Manual, 19K6130						✓
IBM Fibre Channel Problem Determination Guide, 48P9804						✓

FAStT related documents

Table 7 associates each of the following documents related to FAStT operations with its related common user tasks.

Table 7. TotalStorage FAStT related document titles by user tasks

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

How this document is organized

Chapter 1, “Introduction”, on page 1 provides an introduction to the Storage Manager 8.3 product, including information about product resources in addition to this document.

Chapter 2, “Installing storage management station software on AIX systems”, on page 13 provides information about how to install and use Storage Manager 8.3 with an AIX operating system, step-by-step instructions for how to install the AIX software on a management station, and step-by-step instructions for how to install the AIX management software on a host.

Chapter 3, “Installing storage management station software on HP-UX systems”, on page 41 provides information about how to install and use Storage Manager 8.3 with the HP-UX operating system, step-by-step instructions for how to install the HP-UX software on a management station, and step-by-step instructions for how to install the HP-UX management software on a host.

Chapter 4, “Installing storage management station software on Solaris systems”, on page 55 provides information about how to install and use Storage Manager 8.3 with the Solaris operating system, step-by-step instructions for how to install the Solaris software on a management station, and step-by-step instructions for how to install the Solaris management software on a host.

Chapter 5, “Completing the software installation”, on page 73 provides step-by-step instructions for post-installation tasks, such as how to create storage arrays and logical drives, and storage partitioning.

Chapter 6, “Storage Manager 8.3 with high-availability cluster services”, on page 81 provides information about high-availability clustering system options.

Chapter 7, “Copy Services and the FAS*St*T Storage Server”, on page 93 describes FlashCopy, which is a premium feature of the IBM TotalStorage FAS*St*T Storage Manager 8.3. It describes the various components of FlashCopy, followed by a step-by-step instructions for how to use the features.

Appendix A, “AIX system requirements”, on page 139, Appendix B, “HP-UX system requirements”, on page 141, and Appendix C, “Solaris system requirements”, on page 143 list the minimum hardware and software requirements that AIX, HP-UX, and Solaris systems, respectively, must meet to be used in a storage system with Storage Manager 8.3.

Appendix D, “MC/Service Guard configuration details”, on page 145 provides the procedure for how to correct the primary and alternate paths of the imported volume groups that are changed after using **vgimport -m -s** with LVM commands.

Appendix E, “JNI host bus adapter settings”, on page 147 contains JNI host bus adapter (HBA) settings.

Online help

Storage Manager 8.3 provides online help for the Enterprise Management and Subsystem Management windows. These help systems contain information about working with the management domain and about managing storage subsystems.

You can access the help systems from the Enterprise Management and Subsystem Management windows in Storage Manager 8.3. Click **Help** on the toolbar or press F1.

The help systems contain operating information that is common to all operating environments. For operating-system-specific information, see the following chapters and appendices of this document.

- For AIX, see Chapter 2, “Installing storage management station software on AIX systems”, on page 13 and Appendix A, “AIX system requirements”, on page 139.
- For HP-UX, see Chapter 3, “Installing storage management station software on HP-UX systems”, on page 41 and Appendix B, “HP-UX system requirements”, on page 141.
- For Solaris, see Chapter 4, “Installing storage management station software on Solaris systems”, on page 55 and Appendix C, “Solaris system requirements”, on page 143.

Web sites

For the most up-to-date information about IBM FASTT storage servers, go to the following Web site:

www.storage.ibm.com/hardsoft/disk/fastt/

For information about all IBM storage products, go to the following Web site:

www.ibm.com/storage/

For the latest information about operating system and HBA support, clustering support, SAN fabric support, and Storage Manager feature support, see the *TotalStorage FASTT Interoperability Matrix* at the following Web site:

www.storage.ibm.com/disk/fastt/pdf/0217-03.pdf

How to send your comments

Your feedback is important to help us provide the highest quality 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 you are commenting on, such as a page number or table number.

- Mail or fax

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

International Business Machines Corporation
RCF Processing Department
Dept. M86/Bldg. 050-3

5600 Cottle Road
San Jose, CA 95193-0001
U.S.A.

Chapter 1. Introduction

The IBM TotalStorage FAStT Storage Manager Version 8.3 for AIX, HP-UX and Solaris, known as *Storage Manager 8.3*, is a set of client and host tools that allow you to manage IBM FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 storage subsystems from a storage management station.

You can install Storage Manager 8.3 on a *storage management station*, which is the system that is responsible for managing all, or a portion of, a storage network. The storage management station communicates with the network management agents that reside in the managed nodes using a network management protocol, such as Simple Network Management Protocol (SNMP). Storage management commands are sent to the storage subsystem controllers, where the controller firmware validates and runs the commands, and then returns status and configuration information to the client software.

Typically, a storage management station is a remote system, connected to an Ethernet network, that is used to manage one or more storage subsystems. A storage management station can also be a host that is connected to the storage subsystem with a fibre-channel input/output path; you use this same path to manage the attached storage subsystems. Even though you can install the storage management software on a host, the host still uses the Transmission Control Protocol/Internet Protocol (TCP/IP) to communicate with the host-agent. The agent communicates with the controller over the fibre-channel connection through the access volume.

This document provides system administrators with information about how to install, configure and work with Storage Manager 8.3 in AIX, HP-UX and Solaris environments. Before you install Storage Manager 8.3, consult the following documentation:

readme.txt files

Read these first. Text files that contain the latest installation and user information about the storage management software and hardware components are located in each operating system subdirectory on the installation CD. The most recent copies, which supercede this document, are maintained on the following Web site:

www.ibm.com/pc/qtechinfo/MIGR-43839.html

After accepting the license agreement, proceed to the v8.3 link for the appropriate operating system.

IBM FAStT Storage Manager Concepts Guide

Use this reference document to become familiar with the terminology and the features of the Storage Manager 8.3 software. This document is available on the installation CD and maintained at the following Web site:

www.ibm.com/storage/hardsoft/disk/fastt/

For information about how to install Storage Manager 8.3 software on AIX, HP-UX, or Solaris systems, refer to this document. After you have completed the entire installation process, refer to the following online help systems, which contain information that is common to all operating system environments. You can access the help systems from the Enterprise Management and Subsystem Management windows in Storage Manager 8.3 by clicking **Help** on the toolbar or pressing F1.

Enterprise Management Help window

Use this online help system to learn more about working with the entire management domain.

Subsystem Management Help window

Use this online help system to learn more about managing individual storage subsystems.

New system features

Storage Manager 8.3 supports the following:

- FAStT600 Fibre Channel Storage Server
- HACMP, PSSP and GPFS on the FAStT200, FAStT500, FAStT700, and FAStT900 Fibre Channel Storage Servers

Storage Manager 8 premium features

FAStT Storage Manager 8.3 supports the following premium features, which are available for purchase separately from IBM or an IBM Business Partner:

FlashCopy

The FlashCopy premium feature supports the creation and management of FlashCopy logical drives. A FlashCopy is the logical equivalent of a complete physical copy that 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.

Remote Mirror Option

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

This document describes the FlashCopy premium feature. For Remote Mirror Option concepts, configuration procedures, and operation and administration tasks, see *IBM FAStT Remote Mirror Option Installation and User's Guide*.

Storage Partitions

FAStT600

You can purchase four or eight storage partitions for the FAStT600, or you can upgrade from four to eight partitions.

FAStT200, FAStT500, and FAStT700

You can purchase up to 64 storage partitions for the FAStT200, FAStT500 and FAStT700. There is support for up to two host ports in each host and up to eight host ports in each host group.

FAStT900

Storage Partitioning consisting of 16 storage partitions is a standard feature on the FAStT900. You can purchase an upgrade to 64 partitions, which provides support for up to two host ports in each host and up to eight host ports in each host group.

FAStT600 restrictions

The following features are not supported on the FAStT600 Fibre Channel Storage Server:

- HP-UX operating systems
- FlashCopy
- Remote Mirror Option
- HACMP, PSSP and GPFS
- Fibre channel boot

The following restrictions also apply to the FAStT600:

- Supports 2 Gb HBAs only
- Supports 2 Gb link speed only
- Supports FAStT700 EXPs only (maximum of two)
- Supports only one arbitrary loop
- Supports only 42 drives
- Supports a maximum of eight storage partitions

Storage Manager 8 software packages

Storage Manager 8.3 contains the following software packages:

- Storage Manager 8 client software:
 - SMruntime
 - Storage Manager 8 client package (SMclient)
- Storage Manager 8 host software:
 - Storage Manager 8 agent package (SMagent)
 - Storage Manager 8 utility package (SMutil)
 - Storage Manager 8 multipath device drivers (RDAC)

Install the software packages in the sequences shown in Table 8. For installation instructions, see the referenced sections. For an overview of each package, see the subsections that follow the table.

Table 8. Installation sequences of Storage Manager 8.3 software packages by host type

Step	AIX	HP-UX	Solaris
1	SMruntime (see “Installing SMruntime” on page 16)	SMruntime (see “Installing the client software” on page 42)	SMruntime (see “Installing the client software” on page 57)
2	SMclient (see “Installing SMclient” on page 16)	SMclient (see “Installing the client software” on page 42)	SMclient (on Solaris 8 only; see “Installing SMclient” on page 58; supports Solaris 2.6, 2.7, and 2.8)
3	RDAC (see “Installing the RDAC driver” on page 20)	SMagent (see “Installing the SMagent (optional)” on page 44)	SMagent (see “Installing SMagent (optional)” on page 59)
4		SMutil (see “Installing SMutil” on page 45)	SMutil (see “Installing host software on Solaris hosts” on page 59)

Table 8. Installation sequences of Storage Manager 8.3 software packages by host type (continued)

Step	AIX	HP-UX	Solaris
5			RDAC (see “Installing host software on Solaris hosts” on page 59)

Storage Manager 8 client software package

The Storage Manager 8.3 client software, SMclient, is a Java™-based GUI interface that allows you to configure, manage, and troubleshoot FAStT200, FAStT500, FAStT600, FAStT700, or FAStT900 storage servers, and the EXP500 and EXP700 expansion enclosures, in a storage subsystem through a host system or through a storage management station. Specifically, SMclient allows you to:

- Configure disk arrays and logical volumes
- Assign names to arrays and volume groups
- Assign logical volumes to storage partitions
- Replace and rebuild failed disk drives
- Expand the size of arrays
- Expand the size of logical volumes
- Change RAID-type arrays
- Configure and add additional host partition with the premium partitioning feature
- Monitor the status of FAStT storage servers
- Perform troubleshooting and maintenance tasks, such as download firmware to controllers and upgrade nonvolatile static random access memory (NVS RAM)
- View the major events log (MEL)
- Expand the storage capacity using the EXP700 1 Gb or 2 Gb fibre-channel enclosure

Note: FAStT600 supports 2 Gb fibre-channel enclosures only.

- Assign redundant RAID controllers

The SMclient contains two main components:

Enterprise Management

This component allows you to add, remove, and monitor storage subsystems in the management domain.

Subsystem Management

This component allows you to manage the components of an individual storage subsystem.

For more information about this software, see either Chapter 2, “Installing storage management station software on AIX systems”, on page 13, Chapter 3, “Installing storage management station software on HP-UX systems”, on page 41, or Chapter 4, “Installing storage management station software on Solaris systems”, on page 55 and the *IBM FAStT Storage Manager Concepts Guide*.

Storage Manager 8 agent software package

The Storage Manager 8.3 agent (SMagent) package contains the optional host-agent software, which you can use on HP-UX and Solaris host systems to manage storage subsystems through the host fibre-channel connection. The

host-agent software takes requests from a storage management station that is connected to the host through a network connection and passes the requests to the storage subsystem controllers through the fibre-channel I/O path.

For more information about managing storage subsystems through the host agent, see “Host-agent (in-band) management method”.

Storage Manager 8 utility software package

You can use the Storage Manager 8.3 utility (SMutil) package to register and map new logical drives to the operating system. Install SMutil on all HP-UX and Solaris host systems attached to a storage subsystem. The host computers are attached to the storage subsystem through the fibre channel.

RDAC

AIX and Solaris host systems require an RDAC driver for fibre-channel path redundancy. If a FASTT storage server has two controllers, and the operating system does not support multipath I/O, then you can use the RDAC. The RDAC monitors I/O paths; if a component failure occurs in one of the fibre-channel paths, the RDAC reroutes all I/O to another path.

Note: The AIX RDAC driver files are not included on the Storage Manager 8.3 installation CD; you must follow the instructions in “Installing the RDAC driver” on page 20 to download them from the appropriate Web site.

Storage subsystem management methods

The storage management software provides the following two methods for managing storage subsystems:

- The host-agent (in-band) management method. In this method, you manage the storage subsystems through the fibre-channel I/O path to the host.
Restrictions:
 - You cannot use the in-band management method on AIX systems.
 - If both fibre channel connections to the controllers are lost, the SMclient cannot be accessed for problem determination.
- The direct (out-of-band) management method. In this method, you manage the storage subsystems directly over the network through the Ethernet connection to each controller.

Host-agent (in-band) management method

When you use the host-agent (in-band) management method, you manage the storage subsystems through the fibre-channel I/O path to the host. The management information can be processed by the host or passed to the storage management station through the network connection. Figure 2 on page 6 shows the host-agent (in-band) management method.

Restrictions:

- You cannot use the in-band management method on AIX systems.
- If both fibre channel connections to the controllers are lost, the SMclient cannot be accessed for problem determination.




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

Managing storage subsystems using the host-agent (in-band) management method has the following advantages:

- You do not need to run Ethernet cables to the controllers.
- You do not need a Dynamic Host Configuration Protocol (DHCP) bootstrap protocol (BOOTP) server to connect the storage subsystems to the network.
- You do not need to configure the controller network (described in Chapter 2, “Installing storage management station software on AIX systems”, on page 13, Chapter 3, “Installing storage management station software on HP-UX systems”, on page 41, or Chapter 4, “Installing storage management station software on Solaris systems”, on page 55).
- When adding devices, you need to specify a host name or Internet Protocol (IP) address for the host only, not for the individual controllers in a storage subsystem. Storage subsystems that are attached to the host are automatically discovered.

Managing storage subsystems using the host-agent (in-band) management method has the following disadvantages:

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

Important: The access volume uses one of the LUNs. If your host already has the maximum number of LUNs configured, either use the direct-management method or give up a LUN for use as the access volume. For information about your specific configuration, see the appropriate chapter in this document for your operating system environment.

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 name for each controller and attach a cable to the Ethernet ports on each of the storage subsystem controllers. Figure 3 shows the direct (out-of-band) management method.

Restriction: If both network connections to the controllers are lost, the SMclient cannot be accessed for problem determination.




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

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

- The Ethernet connections to the controllers enable a storage management station running SMclient to manage storage subsystems that are connected to a host running one of the following operating systems:
 - AIX

- HP-UX
- Linux
- Microsoft® Windows NT or Windows 2000
- Solaris
- Any other operating systems that are supported by Storage Manager 8.3

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

- It requires two Ethernet cables to connect both storage subsystem controllers to the network.
- When you configure IP addresses for new devices, you must either assign a static IP address or host name for each controller, or place the devices on a network with a DHCP or BOOTP server.
- Other network preparation tasks might be required. For more information, see the system installation guide for your network server.

Operating system requirements

Storage Manager 8.3 supports the following FAStT Fibre Channel Storage Servers in the following environments:

Operating systems	FAStT FC Storage Servers	Where to find more information
IBM AIX versions: 4.3.3 5.1 5.2	FAStT200, FAStT500, FAStT600, FAStT700, FAStT900	Appendix A, "AIX system requirements", on page 139
HP-UX versions: 11.0 (32 bit or 64 bit) 11.11 11i	FAStT200, FAStT500, FAStT700, FAStT900	Appendix B, "HP-UX system requirements", on page 141
Sun Solaris versions: 2.6 2.7 2.8	FAStT200, FAStT500, FAStT600, FAStT700, FAStT900	Appendix C, "Solaris system requirements", on page 143

FAStT600:

HP-UX operating systems are not supported on the FAStT600 Fibre Channel Storage Server.

Setting up IP addresses for FAStT storage controllers

To use the out-of-band management method without setting up a DHCP BOOTP server, you must assign IP addresses to the FAStT controllers using command-line interface (CLI) commands through serial cables that are connected to a terminal emulator.

Before you begin: Contact your network administrator to obtain the IP address and associated host name for each controller in every storage subsystem on the network, and make a note of those values for reference.

- Ensure that you have the following addresses:

- An IP address for each controller
- A submask address
- Download the latest version of HyperTerminal Private Edition (6.3 or higher) from the following Web site:
www.hilgraeve.com

Complete the following steps to set up the FAStT controller IP addresses using serial ports:

1. Stop all I/O to the FAStT controllers.
2. Connect a null modem serial cable from one of the controllers to a system with a terminal emulator available.
3. Open the HyperTerminal and from the menu bar click **File** → **Properties** → **Configure**. Choose the following settings:
 - 115200 Baud
 - 8 Data Bits
 - 1 Stop Bit
 - No parity
 - XON/XOFF Flow Control
4. Connect to the FAStT storage server and send a break signal (Ctrl+Break for most emulators).
5. Repeat this step until the following message is displayed:

Press the space bar for baud rate within 5 seconds.

6. Press the space bar to ensure the correct baud rate setting.
7. Send another break signal; the following message is displayed:

Press within 5 seconds: ESC for SHELL, BREAK for baud rate.

8. Press Escape to access the shell of the controller.
9. Type the following password: `infiniti`.
10. Type `netCfgShow` to show the current network configuration.

Note: The default following IP address settings are set by manufacturing:

- Controller A = 192.168.10.101
- Controller B = 192.168.10.102
- IP address mask = 255.255.255.0

11. Type `netCfgSet` to change the network configuration information.

Note: Press Enter to advance to the next field. Type the new IP address in the **My IP Address** field.

12. Assign an IP address to the controller.
13. Type `netCfgShow` to verify the new network settings.
14. Disconnect from the first controller and connect to the second controller.
15. Repeat steps 1 - 14 to assign the second IP address to the second controller.
16. Turn the controller unit off and on to restart the FAStT storage server.

Enabling Storage Partitioning

The following Storage Partition Premium Features are available for purchase separately from IBM or an IBM Business Partner:

FAStT200, FAStT500, FAStT700

You can purchase up to 64 storage partitions.

FAStT900

Storage Partitioning consisting of 16 storage partitions is a standard feature of the FAStT900. You can purchase an upgrade from 16 to 64 partitions.

FAStT600

You can purchase four partitions, eight partitions, or an upgrade from four to eight partitions.

To enable the Storage Partitioning Premium Feature, perform the procedures that are described in this section in the following order:

1. Obtaining the feature enable identifier
2. Generating the feature key file and enabling Storage Partitioning

Before you begin: Gather the following information:

- Your Feature Activation Code, which accompanies the Storage Partitioning Premium Feature
- Your controller unit IBM serial number, which is printed on the label on the back of your FAStT controller unit

Ensure that your controller unit and expansion units are connected, powered on and configured.

Obtaining the feature enable identifier

Each storage subsystem has its own unique feature enable identifier. This identifier ensures that a particular feature key file is applicable only to that storage subsystem. Complete the following steps to obtain the feature enable identifier:

1. Click **Start** → **Programs** → **FAStT Storage Manager 8.3 Client**. The Enterprise Management window opens.
2. In the left-hand pane, double-click the storage subsystem for which you want to enable Storage Partitioning or a Storage Partitioning upgrade. The Subsystem Management window opens for the selected storage subsystem.
3. In the Subsystem Management window, click **Storage Subsystem** → **Premium Features** → **List**. The Premium Feature List window opens and displays the feature enable identifier.
4. Record the feature enable identifier.
5. Close the Premium Feature List window.

Generating the feature key file and enabling Storage Partitioning

You enable the Storage Partitioning feature using a tool that is located at the following Web site:

www.ibm.com/storage/fasttkeys

1. Generate and store the key activation file by completing Step 1 on the Web site.
2. Download the key activation file to your controller unit by completing Step 2 on the Web site.
3. Enable Storage Partitioning by completing Step 3 on the Web site.

| 4. Verify that Storage Partitioning was successfully enabled by completing Step 4
| on the Web site.

| Your system is now ready to use Storage Partitioning or the Storage Partitioning
| upgrade.

| If Storage Partitioning becomes disabled, you can access the Web site and repeat
| this process. For any additional assistance, please contact your local IBM service
| provider.

Chapter 2. Installing storage management station software on AIX systems

This chapter provides the following specific information for AIX operating systems:

- Hardware and firmware requirements
- AIX restrictions
- Installing client software
- Performing the initial configuration
- Creating direct-attached and SAN-attached configurations
- Installing RDAC
- Implementing dynamic volume expansion (DVE) and dynamic capacity expansion (DCE)
- Common disk array errors

Hardware and firmware requirements

Table 9 lists the FAStT storage servers that are supported with Storage Manager 8.3.

Table 9. FAStT storage servers that are supported on AIX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT900 RAID Controller Enclosure Unit	1742-90U	Appware 05.30.09.00 NVSRAM N1742F900R830V04
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 05.30.09.00 NVSRAM N1742F700R830V03
IBM TotalStorage FAStT600 RAID Controller Enclosure Unit	1722-60U, 1722-60X	Snapware 08.33.05.00 NVSRAM N1722F600R833V02
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 05.30.09.00 NVSRAM N3552F500R830V03
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05

Table 10 lists the supported versions of hardware to use with Storage Manager 8.3.

FAStT600: The FAStT600 does not support all of the expansion units, switches and HBAs that are listed in Table 10. Please note the checkmarks that indicate which versions are supported on the FAStT600.

Table 10. Supported versions of hardware for AIX systems

Product Name	Model	Product release and firmware version	Supported on FAStT600
IBM TotalStorage FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9321	✓
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9166	

Table 10. Supported versions of hardware for AIX systems (continued)

Product Name	Model	Product release and firmware version	Supported on FAST600
IBM HA Emulex LP7000	FC 6227	S1F3.22A1	
IBM HA Emulex LP9000	FC 6228	C1D3.82A1	✓
Brocade switch	2109-S08 2109-S16	2.6.0j	
Brocade switch	2109-F16	3.0.2k	✓
Brocade switch	2109-M12, 2109-F32	4.0.2b	✓
Brocade switch	2109-F08	3.0.2k	✓
McData switch	2032-064, 2032-140	4.1	✓
McData switch	2031-016	4.1	
McData switch	2031-216, 2031-224	4.1	✓
InRange switch	2042-001	4.1.05	✓
Cisco switch	2062-D01, 2062-D07, 2062-T07	GA level	✓

Note: Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. To ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers, read the readme.txt file that is shipped with the product. You can also find a copy of the readme.txt file at the following Web site:

www.ibm.com/pc/qtechinfo/MIGR-50177.html

After accepting the license agreement, proceed to the v8.3 link for the appropriate operating system.

You can download the latest firmware, NVSRAM and device driver versions at the following Web sites:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

AIX restrictions

The following restrictions apply to FAST200, FAST500, FAST600, FAST700, and FAST900 storage servers:

- F-RAID Manager is not supported.
- The maximum number of partitions per AIX host, per FAST storage server, is two.

- Each AIX host can support two or four host bus adapters and up to two FAStT storage partitions, each requiring two HBA connections per FAStT storage server. Additional HBA pairs can be added to support additional FAStT storage servers.
- Direct-attach configurations are restricted to single-initiator configurations only. Only one connection to each FAStT mini-hub is allowed.
- Single HBA configurations are allowed, but each single HBA configuration requires that both controllers in the FAStT be connected to a switch, and within the same SAN zone as the HBA.
- Single-switch configurations are allowed, but each HBA and FAStT controller combination must be in a separate SAN zone.
- All volumes that are configured for AIX must be mapped to an AIX host group. Connecting and configuring to volumes in the default host group is not allowed. See “Performing the initial configuration of storage subsystems on AIX hosts” on page 16.
- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.
- Booting from a FAStT device is supported only on AIX 5.1 and 5.2.
- When you boot from a FAStT device, both paths to the boot device must be up and operational.
- Path failover is not supported during the AIX boot process. Once the AIX host has booted, failover operates normally.
- You cannot perform a SAN boot from the FAStT600 or FAStT900 storage servers.

Installing the client software

Use the procedures in this section to install the client software on an AIX storage management station.

Install the software in the following order:

1. SMruntime
2. SMclient

Note: The SMclient is dependent on SMruntime, which is a Java compiler for the SMclient and must be installed first.

Prerequisites

Before installing the software, ensure that the following conditions are met:

- The AIX host on which you are installing SMruntime meets the minimum hardware and software requirements described in “Hardware and firmware requirements” on page 13 and “Software requirements” on page 139.
- The correct filesets are present on the system.

Note: If the filesets are not present, follow the instructions in “Installing the RDAC driver” on page 20 to download them from the appropriate Web site.

- For a list of AIX 4.3.3 filesets, see Table 20 on page 139.
- For a list of AIX 5.1 filesets, see Table 21 on page 139.
- For a list of AIX 5.2 filesets, see Table 22 on page 140.

Installing SMruntime

Adjust these instructions as required for your specific installation. No restart is required during the installation process.

1. Install SMruntime by typing the following command:

```
# installp -a -d /complete path name  
/SMruntime.aix-08.30.65.00.bff SMruntime.aix.rte
```

2. Verify that the installation was successful by typing the following command:

```
# ls1pp -ah SMruntime.aix.rte
```

The verification process returns a table that describes the software installation, including the install package file name, version number, action and action status. If the verification process returns an error, contact your service representative.

Installing SMclient

In the following procedure, the installation CD is mounted at /dev/cdrom. Adjust these instructions as required for your specific installation. No restart is required during the client installation process.

1. Install SMclient by typing the following command:

```
# installp -a -d /complete path name/  
SMclient.aix-08..33.G5.03.bff SMclient.aix.rte
```

2. Verify that the installation was successful by typing the following command:

```
# ls1pp -ah SMclient.aix.rte
```

The verification process returns a table that describes the software installation, including the install package file name, version number, action and action status. If the verification process returns an error, contact your support representative.

Performing the initial configuration of storage subsystems on AIX hosts

If you have purchased a Storage Partitioning Premium Feature, ensure that the feature is enabled. See “Enabling Storage Partitioning” on page 10 for more information.

Complete the following procedure to configure Storage Manager 8.3 for an AIX system. This can be done from the SMclient running on either an AIX or a non-AIX system.

1. To set up the storage subsystem for AIX and the AIX SMclient, the subsystem must be physically configured for direct management through the Ethernet connections on each controller. Install SMclient before configuring the subsystem.

Note: See “Setting up IP addresses for FASiT storage controllers” on page 8 for information about assigning IP addresses to the controllers.

2. After the disk subsystem is configured on the network, start the SMclient on the host server by typing the following command:

```
# /usr/SMclient/SMclient
```

3. Complete the following steps to specify the IP addresses of the controllers:
 - a. In the Enterprise Management window, click **Edit** → **Add Device**.
 - b. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.
 - c. Type the IP address of the second controller and click **Add**, and then click **Close**.

The storage subsystem is shown as a direct network attachment. Double-click **Storage Subsystem** to open the Subsystem Management window.

Updating FASSt firmware and NVSRAM

Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. Read the readme.txt file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

Perform the following steps to update the FASSt controller firmware and NVSRAM:

1. In the Subsystem Management window, click **View** → **Storage Subsystem Profile** and review the summary portion of the output. Verify that the controller firmware and NVSRAM are at the correct versions. If they are not at the correct versions, upgrade the firmware, and then upgrade the NVSRAM.

Note: Even though concurrent firmware upgrades are supported, I/O should be quiesced before you upgrade the firmware. Concurrent upgrades of NVSRAM, however, are not supported.

2. Complete the following steps to upgrade the controller firmware:
 - a. Ensure that all I/O to the controllers has been stopped.
 - b. Download the correct version of the firmware from the Web site. Place the file in a designated directory on the host system.
 - c. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **Firmware**.
 - d. Type or select the full pathname of the firmware file.
 - e. Click **OK** to update both controllers.
 - f. Click **Yes** to start the download. A new window opens.
 - g. Close the Subsystem Management window and then reopen it to complete the firmware update.
 - h. Verify that the firmware was successfully installed.
3. Complete the following steps to upgrade the NVSRAM:
 - a. Ensure that all I/O to the controllers has been stopped.
 - b. Download the correct version of NVSRAM from the Web site. Place the file in a designated directory on the host system.
 - c. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **NVSRAM**.

- d. Type or select the full pathname of the NVSRAM directory.
- e. Double-click the NVSRAM filename or click **OK** to select the correct file.
- f. Click **Yes** to start the download.
- g. Verify that the NVSRAM was successfully installed.

Setting up an AIX host group

Perform the following steps to set up an AIX host group:

1. Click the **Mappings View** tab on the Subsystem Management window.
2. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.
3. Type the name of the new host group (for example, AIX). Click **Add**, and then click **Close**.
4. Highlight the new host group and click **Mappings** → **Define** → **Host**.
5. Define the new host. Type the name of the AIX host to which the storage subsystem is attached.
 - a. Click **Add**, and then click **Close**.
 - b. Highlight the host that you just added and right-click **Define New Host Port**.
 - c. Select the desired host port for the first HBA, and then change the host type to **AIX** and click **Add**.

Note: Failure to change the host type from the default to AIX will cause undesired results

- d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

After you set up an AIX host group, perform the following steps to verify that the host ports match the AIX host:

1. Type the following command:

```
lsdev -Cc adapter | grep fcs
```

A list that contains all the HBAs that are in the system is displayed, as shown in the following example:

```
# lsdev -Cc adapter |grep fcs
fcs0    Available 20-58    FC Adapter
fcs1    Available 30-70    FC Adapter
fcs2    Available 40-60    FC Adapter
fcs3    Available 90-58    FC Adapter
fcs4    Available B0-70    FC Adapter
fcs5    Available C0-60    FC Adapter
```

2. Identify the fcs number of the HBA that is connected to the FAS*t*T.
3. Type the following command:

```
lscfg -v1 fcs# |grep Network
```

where *fcs#* is the fcs number of the HBA that is connected to the FAS*t*T.

The network address number of the HBA is displayed, as in the following example:

```
# lscfg -vl fcs0|grep Network
Network Address.....10000000C925514C
```

Note: A network address is also known as a worldwide port name (WWPN).

4. Verify that the network address number matches the host port number that displays in host partition table of the FASTT SMclient.
5. Repeat this procedure to verify the second host port.

Creating direct and SAN-attached configurations

Storage Manager 8.3 supports FASTT200, FASTT500, FASTT600, FASTT700, and FASTT900 storage servers in direct-attached AIX configurations or in a SAN environment through switches in AIX configurations.

FASTT600:

For the FASTT600, you must set the link speed at 2 Gb on the controllers and on the FASTT700 EXP.

Creating a direct-attached configuration

To create a direct-attached configuration, you must ensure that:

- One or two AIX servers can be connected to the FASTT storage server:
 - FASTT200 can support one AIX server.
 - FASTT500, FASTT600, FASTT700, and FASTT900 can support two AIX servers.
- Two-server FASTT500, FASTT700 or FASTT900 configurations require four host-side mini-hubs, each with exactly one fibre-channel connection from each HBA to a mini-hub.
- There are two or four HBAs per FASTT storage server. Each pair must be configured to one FASTT partition.
- No external hubs are being used.

To set up a direct-attached configuration, follow these steps:

1. Connect the HBAs to each controller or mini-hub port of the FASTT storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

To create a SAN-attached configuration, you must ensure that:

- Multiple fibre-channel HBAs within the same server cannot “see” the same FASTT controller port.
- The IBM fibre-channel HBAs are isolated from each other if they are connected to the same switch that is connected to the same FASTT controller port.
- Each fibre-channel HBA and controller port must be in its own fabric zone, if they are connecting through a single fibre-channel switch, such as 2109-F16.

See the documentation that is provided by the switch manufacturer for more information about zoning. Multiple FASTT devices can be configured to the same set of fibre-channel HBAs through a fibre-channel switch.

To set up a SAN-attached configuration, follow these steps:

1. Connect the HBAs to the switch or switches.
2. Connect the FAStT storage subsystems to the switch or switches.
3. Set the required zones on the fibre-channel switch or switches, if applicable.

Note: For information about zoning and enabling zones, see the documentation that is provided by the switch manufacturer.

4. Configure and verify the configuration.

FAStT600:

The FAStT600 Fibre Channel Storage Server only supports the use of 2 Gb fibre-channel switches. Set the port speed on the 2 Gb fibre-channel switches to 2 Gb.

Installing the RDAC driver

After you install the client software and configure your storage subsystems, use these instructions to install the appropriate Storage Manager 8.3 RDAC device driver.

All AIX hosts in your storage subsystem must have the RDAC multipath driver installed. This section describes how to check the current RDAC driver program driver version level, update the RDAC device driver, and verify that the RDAC update is complete.

Prerequisites

This section lists the prerequisites for installing the RDAC driver version levels on an AIX system and describes how to verify the current RDAC drive driver version level.

Prerequisites for installing RDAC

See Table 20 on page 139, Table 21 on page 139, or Table 22 on page 140 for a list of required driver version level filesets.

Verifying RDAC software

Use the **lslpp** command to verify that the correct driver versions are installed:

```
# lslpp -ah filename
```

Where *filename* is one of the required filesets listed in Table 20 on page 139, Table 21 on page 139 or Table 22 on page 140.

For example:

```
# lslpp -ah devices.fcp.disk.array.rte
```

Installing the RDAC driver

Complete the following procedure to update the RDAC driver version (devices.fcp.disk.array.rte) on an AIX system. Repeat this procedure for all AIX systems that are connected to the storage subsystem. You do not need to perform this installation if you have verified that the RDAC level is correct.

1. Go to one of the following Web sites:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

2. Click **Downloads** in the downloads section of the Web page.
3. Scroll to the operating system-specific updates section of the Web page.
4. Follow the link to the appropriate sets of files for your operating system and follow the installation instructions.
5. Verify that the correct version of the software was successfully installed by typing the following command:

```
# ls1pp -ah devices.fcp.disk.array.rte
```

The verification process returns a table that describes the software installation, including the installation package fileset name, version number, action, and action status. If the verification process returns an error, contact your customer service representative. If it does not return an error, then you are finished installing the updated RDAC driver on this AIX system.

6. Configure the devices for the software changes to take effect by typing the following command:

```
# cfmgr -v
```

AIX configuration information

This section contains the following AIX configuration information:

- Viewing and setting attributes of the RDAC driver for AIX
- Verifying the installation and configuration of AIX hosts
- Identifying the controller ID numbers
- Identifying device names and bus numbers
- Identifying logical drives by operating system device names
- Identifying FlashCopy volume types

Viewing and setting attributes of the RDAC driver for AIX

The RDAC driver must be installed on all AIX hosts that will be attached to a FAS*T*T storage subsystem. The RDAC driver creates the following devices that represent the FAS*T*T storage subsystem configuration:

dar The disk array router represents the entire array, including the current and the deferred paths to all LUNs (hdisk*s* on AIX).

dac The disk array controller devices represent a controller within the storage subsystem. There are two dacs in the storage subsystem.

hdisk These devices represent individual LUNs on the array.

When these devices are configured, the Object Data Manager (ODM) is updated with default parameters. In most cases and for most configurations, the default parameters are satisfactory. However, there are some parameters that can be modified for maximum performance and availability. See “hdisk attribute settings” on page 28.

See “Viewing Object Data Manager (ODM) attributes in AIX” on page 26 for information about using the **lsattr** command to view attribute settings on an AIX system.

Definitions of attribute settings

Some attributes for dars, dacs, and hdisks are for information purposes only, showing how the FAStT is configured or its current state. You can modify other attributes with SMIT or with the AIX **chdev -p** command. Attributes with True in the Changeable column can be modified from their default settings. False means that they are for informational or state purposes only. Another method of determining attributes that can be modified is with the **lsattr -el** command. Attributes with True in the last column of the output can be modified. See the examples in Table 11, Table 12 on page 23, and Table 13 on page 24. Display the default with the **lsattr -dl** command.

Table 11. Attributes for dar devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>act_controller</i>	List of controllers in the active state at the time of configuration.	False	Set at configuration time by the RDAC software.
<i>all_controller</i>	List of controllers that comprise this array; usually there are two dac devices.	False	Set at configuration time by the RDAC software.
<i>held_in_reset</i>	Name of the controller that was in the held-in-reset state at the time of configuration, or none if no controllers were in that state.	True	Set at configuration time by the RDAC software. Should not be changed.
<i>load_balancing</i>	Indicator that shows whether load balancing is enabled (yes) or disabled (no); see the definition of the <i>balance_freq</i> attribute for more information.	True	Yes or No. See restrictions on use.
<i>autorecovery</i>	Indicator that shows whether the device returns the array to dual-active mode when it detects proper operation of both paths and controllers (yes) or not (no).	True	yes or no (see restrictions on use)
<i>lun_bitmap</i>	Bitmap that shows which controller owns which LUN at the time of configuration.	True	While changeable, this should not be changed. LUN ownership should be controlled by the Storage Manager.

Table 11. Attributes for dar devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>hlthchk_freq</i>	Number that specifies how often health checks are performed, in seconds.	True	1 - 9999. Should not be changed
<i>aen_freq</i>	Number that specifies how often polled AEN checks are performed, in seconds.	True	1 - 9999. Should not be changed
<i>balance_freq</i>	If <i>load_balancing</i> is enabled, number that specifies how often the system performs load-balancing on the array, in seconds.	True1 - 9999. Should not be changed	1 - 9999 - should not be changed
<i>fast_write_ok</i>	Indicator that shows whether fast-write write-caching is available for this system (yes) or not (no).	False	Yes or No. State of FAStT configuration.
<i>cache_size</i>	Cache size for both controllers, in megabytes; 0 if the sizes do not match.	False	512 or 1024. Set by FAStT.
<i>switch_retries</i>	Number that specifies how many times to retry failed switches, in integers.	True	0 - 255. Should not be changed.

Table 12. Attributes for dac devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>passive_control</i>	Indicator that shows whether this controller was in passive state at the time of configuration (yes) or not (no).	False	Yes or No. State of FAStT configuration.
<i>alt_held_reset</i>	Indicator that shows whether the alternate controller was in the held-in-reset state at the time of configuration (yes) or not (no).	False	Yes or No. State of FAStT configuration.
<i>controller_SN</i>	Serial number of this controller.	False	Set by FAStT.

Table 12. Attributes for dac devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>ctrl_type</i>	Type of array this controller belongs to. A value of 3542 indicates FAStT200; a value of 3552 indicates FAStT500; a value of 1722 indicates FAStT600; a value of 1742 indicates FAStT700; a value of 1742-900 indicates FAStT900.	False	3542, 3552, 1742, 1722, 1742-900. Set by FAStT.
<i>cache_size</i>	Cache size of this controller, in megabytes.	False	512, 1024. Set by FAStT.
<i>scsi_id</i>	SCSI identifier of this controller.	False	Set by SAN, reported by AIX.
<i>lun_id</i>	Logical unit number of this controller.	False	Set by FAStT.
<i>utm_lun_id</i>	Logical unit number of this controller, or none if UTM (access volumes) is not enabled.	False	0 - 31. Set by FAStT Storage Manager.
<i>location</i>	User-defined location label for this controller; the system does not use this value.	True	Set by FAStT Storage Manager.
<i>ww_name</i>	Fibre-channel worldwide name of this controller.	False	Set by FAStT.
<i>GLM_type</i>	GLM type used for this controller.	False	High or Low. Set by FAStT.

Table 13. Attributes for hdisk devices

Attribute	Definition	Changeable (T/F)	Possible value
<i>pvid</i>	AIX physical volume identifier, or none if not set.	False	Set by AIX.
<i>q_type</i>	Queueing type for this device; must be set to simple .	False	Set by AIX. Must be "simple".
<i>queue_depth</i>	Number that specifies the depth of the queue based on system configuration; reduce this number if the array is returning a BUSY status on a consistent basis.	True	1 - 64

Table 13. Attributes for *hdisk* devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>reserve_lock</i>	Indicator that shows whether the <i>fcarray</i> issues a SCSI Remove command every time a device is opened or when a Test Unit Ready sequence is issued by the driver (yes) or not (no).	True	Yes or No.
<i>write_cache</i>	Indicator that shows whether write-caching is enabled on this device (yes) or not (no); see the definition of the <i>cache_method</i> attribute for more information.	True	Yes or No.
<i>size</i>	Size of this volume.	False	Set by FAST.
<i>raid_level</i>	Number that specifies the RAID level of this device.	False	0, 1, 3, 5. Set by FAST Storage Manager.
<i>rw_timeout</i>	Number that specifies the read/write timeout value for each read/write command to this array, in seconds; usually set to 30.	True	30 - 180. Should not be changed from default.
<i>reassign_to</i>	Number that specifies the timeout value for FC reassign operations, in seconds; usually set to 120.	True	0 - 1000. Should not be changed from default.
<i>scsi_id</i>	SCSI identifier at the time of configuration.	False	Set by SAN, reported by AIX.
<i>lun_id</i>	Logical unit number of this device.	False	0 - 31. Set by FAST Storage Manager.

Table 13. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<i>cache_method</i>	<p>If <i>write_cache</i> is enabled, the write-caching method of this array; set to one of the following:</p> <ul style="list-style-type: none"> • default. Default mode; the word "default" is not seen if <i>write_cache</i> is set to yes. • fast_write. Fast-write (battery-backed, mirrored write-cache) mode. • fw_unavail. Fast-write mode was specified but could not be enabled; write-caching is not in use. • fast_load. Fast-load (non-battery-backed, non-mirrored write-cache) mode. • fl_unavail. Fast-load mode was specified but could not be enabled. 	True	Default, fast_write, fast_load, fw_unavail, fl_unavail.
<i>prefetch_mult</i>	Number of blocks to be prefetched into read cache for each block read.	True	0 - 100.
<i>ieee_volname</i>	IEEE unique volume name identifier for this volume.	False	Set by FAStT.

Viewing Object Data Manager (ODM) attributes in AIX

To view the attributes for dars, dacs, and hdisks, use the **lsattr** command. The following examples demonstrate how to use the **lsattr -EI** command, which shows the attributes that are currently set on the system. To view the default settings, use the **-DI** flag instead of **-EI**.

This example shows the attribute settings for a disk array router (dar0):

```

# lsattr -El dar0
act_controller dac0,dac1           Active Controllers           False
all_controller dac0,dac1           Available Controllers        False
held_in_reset  none                 Held-in-reset controller    True
load_balancing no                     Dynamic Load Balancing       True
autorecovery   no                     Autorecover after failure is corrected True
lun_bitmap     00000000000000000000000000001010 LUN Ownership bitmap        True
hlthchk_freq   600                     Health check frequency in seconds True
aen_freq       600                     Polled AE frequency in seconds True
balance_freq   600                     Dynamic Load Balancing frequency in seconds True
fast_write_ok  yes                     Fast Write available         False
cache_size     512                     Cache size for both controllers False
switch_retries 5                     Number of times to retry failed switches True

```

This example shows the attribute settings for a disk array controller (dac0):

```

# lsattr -El dac0
passive_control no                 Passive controller           False
alt_held_reset no                 Alternate held in reset      False
controller_SN  1T04611221                Controller serial number     False
ctrl_type      3552                  Controller Type              False
cache_size     512                     Cache Size in MBytes        False
scsi_id        0x210013                SCSI ID                      False
lun_id         0x0                     Logical Unit Number         False
utm_lun_id     none                    Logical Unit Number         False
location       none                    Location Label               True
ww_name        0x200800a0b80c1948 World Wide Name              False
GLM_type       low                     GLM type                     False

```

This example shows the attribute settings for a LUN on the array (hdisk4):

```

# lsattr -El hdisk4
pvid           none                     Physical volume identifier    False
q_type         simple                    Queuing Type                  False
queue_depth    30                       Queue Depth                   True
reserve_lock   yes                        RESERVE device on open       True
write_cache    yes                        Write Caching enabled         True
size           69270                     Size in Mbytes                False
raid_level     5                        RAID Level                    False
rw_timeout     30                       Read/Write Timeout value     True
reassign_to    120                       Reassign Timeout value       True
scsi_id        0x210013                SCSI ID                       False
lun_id         0x0000000000000000          Logical Unit Number          False
cache_method   fast_write                    Write Caching method         True
prefetch_mult  0                          Multiple of blocks to prefetch on read True
ieee_volname   600A0B80000C1A70000000013B45FEA8 IEEE Unique volume name      False

```

Changing ODM attribute settings in AIX

As discussed earlier, ODM attributes for the RDAC driver and FASTt can be modified. This section shows the settings that can be used for best performance and availability.

hdisk attribute settings

This section lists the attribute settings that you should use for hdisk devices and shows how to set them using the **chdev -l** command. To make the attribute changes permanent in the Customized Devices object class, use the **-P** option.

Some attributes can be changed from both the SMclient and through AIX. To make these changes permanent, the modifications must be made by changing the AIX ODM attribute. The specific hdisk attributes affected are:

- cache_write
- cache_method
- prefetch_mult

If the changes are made through the SMclient, they will operate properly until you either restart the host or restart cfgmgr. To make the changes permanent, you must use SMIT or the **chdev -P** command.

Attention: Controller cache mirroring should not be disabled while cache_write is enabled. If this condition exists, the RDAC software will automatically re-enable it the next time the system is restarted, or when **cfgmgr** is run.

Restrictions for the load_balancing attribute: The *load_balancing* attribute should only be set to **yes** in single-host configurations.

Setting the queue depth for hdisk devices: Setting the *queue_depth* attribute to the appropriate value is important for system performance. For large FASTT configurations with many volumes and hosts attached, this is a critical setting for high availability. If you set this attribute incorrectly, it could result in the loss of filesystems and system panics.

Use the following formula to determine the maximum queue depth for your system:

$$512 / (\text{number-of-hosts} * \text{LUNs-per-host})$$

For example, a system with four hosts, each with 32 LUNs (the maximum number of LUNs per AIX host), would have a maximum queue depth of 4:

$$512 / (4 * 32) = 4$$

In this case, you would set the *queue_depth* attribute for hdiskX as follows:

```
# chdev -l hdiskX -a queue_depth=4 -P
```

Attention: If you do not set the queue depth to the proper level, you might experience loss of filesystems and system panics.

Verifying the installation and configuration of AIX hosts

After you set the ODM attributes, FASTT volumes, and configurations, and before you mount your file systems and install your applications, use the following information to verify that all of your FASTT device names and paths are correct and that AIX recognizes your dars, dacs, and hdisks.

Initial device identification

After the FASTT storage subsystem has been set up, volumes have been assigned to the host, and the RDAC driver has been installed, type the following command to probe for the new devices:

```
# cfgmgr -v
```

Next, use the **lsdev -Cc disk** command to see if the RDAC software recognizes each FAStT volume as a “3542 (200) Disk Array Device”, each FAStT500 volume as a “3552 (500) Disk Array Device”, each FAStT600 volume as a “1722 (600) Disk Array Device”, each FAStT700 volume as a “1742 (700) Disk Array Device”, or each FAStT900 volume as a “1742-900 Disk Array Device”. The following example shows the results of the command for a set of FAStT500 LUNs:

```
# lsdev -Cc disk
hdisk0 Available 10-88-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk32 Available 31-08-01 3552 (500) Disk Array Device
hdisk33 Available 91-08-01 3552 (500) Disk Array Device
hdisk34 Available 31-08-01 3552 (500) Disk Array Device
hdisk35 Available 91-08-01 3552 (500) Disk Array Device
```

Identifying the controller ID numbers


One of the attributes listed by the **lsattr -E1** command is the controller serial number (*controller_SN*) of that dac.

Perform the following steps to display the FAStT Controller Properties window:

1. In the Physical view of the Subsystem Management window, select a controller.
2. Right-click **Controller** → **Properties**.

Figure 4 on page 30 shows the Controller Properties window of the controller in slot A. This controller has a serial number of 1T04810361 and is represented by dac1.

```
# lsattr -E1 dac1
passive_control no Passive controller False
alt_held_reset no Alternate held in reset False
controller_SN 1T14610048 Controller serial number False
ctrl_type 3552 Controller Type False
cache_size 512 Cache Size in MBytes False
scsi_id 0x210213 SCSI ID False
lun_id 0x0 Logical Unit Number False
utm_lun_id none Logical Unit Number False
location Location Label True
ww_name 0x200600a0b80c213d World Wide Name False
GLM_type low GLM type False
```



SJ000709

Figure 4. Controller Properties window

Identifying device names and bus numbers

After the operating system device names are found, those names must be correlated to the preferred and failover paths of the FASTt device, and then from each path to its associated logical drive. As mentioned earlier in this chapter, the RDAC software uses `dars`, `dacs`, and `hdisks` to represent the FASTt storage subsystem. `dars` represent the disk array routers; `dacs` represent the disk array controllers, and `hdisks` represent the logical drives, or FASTt volumes.

AIX provides the following commands to help you determine the FASTt configuration, and to get information about device names and bus numbers:

lsdev

Displays devices and their characteristics. The **lsdev** command shows the state of the devices at startup time, or the last time that the **cfgmgr -v** command was run.

lsattr

Displays device attributes and possible values. Attributes are only updated at startup time, or the last time that the **cfgmgr -v** command was run.

fget_config

Displays controllers and `hdisks` that are associated with a specified FASTt (`dar`).

The **fget_config** command shows the current state and volume (hdisk) ownership.

There are several ways to correlate a system's configuration and monitor the state of FAStT storage subsystems. The use of commands in validating configurations is illustrated in Figure 5.




Figure 5. Use of commands applied to validating system configuration

Using the **lsdev** command

The following example uses the **lsdev** command to show the status of the dar, which represents a FAStT storage subsystem. This example shows dar as a machine type 3542, which is a FAStT200. It is in the Available state, which is the state at the time when the device was last configured by AIX.

```
# lsdev -C |grep dar0
dar0      Available          3542      (200) Disk Array Router
```

The following example uses the **lsdev** command to show the status of two dacs, which represent the FAStT storage subsystem controllers. As in the previous example, this shows a FAStT200 in an Available state. The third column shows the location code. In this example, each dac has its own location or *path*, which are represented by the values 11-08-01 and 14-08-01. Each AIX system has its own set of location codes that describe the internal path of that device, including bus and host-adapter locations. See the service manual for your system type to identify device locations.

```
# lsdev -C |grep dac
dac0      Available 11-08-01      3542      (200) Disk Array Controller
dac1      Available 14-08-01      3542      (200) Disk Array Controller
```

The following example uses the **lsdev** command to show the status and location codes of eight FAStT200 hdisks. Notice that the location codes of the odd hdisks (1, 3, 5, 7) correspond to the same location code of dac0 in the previous example, and that the even hdisks (2, 4, 6, 8) correspond to the same location code of dac1. This means that the preferred paths for I/O for hdisk1, hdisk3, hdisk5, and hdisk7 is through dac0. The failover path would be through dac1. Conversely, the preferred path for the even-numbered hdisks would be through dac1, and failover path through dac0.

```
# lsdev -Cc |grep hdisk
hdisk0    Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1    Available 11-08-01      3542      (200) Disk Array Device
hdisk2    Available 14-08-01      3542      (200) Disk Array Device
hdisk3    Available 11-08-01      3542      (200) Disk Array Device
hdisk4    Available 14-08-01      3542      (200) Disk Array Device
hdisk5    Available 11-08-01      3542      (200) Disk Array Device
hdisk6    Available 14-08-01      3542      (200) Disk Array Device
hdisk7    Available 11-08-01      3542      (200) Disk Array Device
hdisk8    Available 14-08-01      3542      (200) Disk Array Device
```

You can determine the full preferred path from hdisk to fibre-channel fabric by searching for the first two numbers of the location code using the **grep** command, as in the following example:

```
# lsdev -C|grep 11-08
fcs0      Available 11-08          FC Adapter
fscsi0    Available 11-08-01      FC SCSI I/O Controller Protocol Device
dac0      Available 11-08-01      3542      (200) Disk Array Controller
hdisk1    Available 11-08-01      3542      (200) Disk Array Device
hdisk3    Available 11-08-01      3542      (200) Disk Array Device
hdisk5    Available 11-08-01      3542      (200) Disk Array Device
hdisk7    Available 11-08-01      3542      (200) Disk Array Device
hdisk8    Available 11-08-01      3542      (200) Disk Array Device
```

Using the **fget_config** command

The **fget_config** command displays the state of each controller in a FAStT array, and the current path that is being used for I/O for each hdisk. The following example shows that both controllers (dac0 and dac1) are in the Active state. This is normal when the FAStT storage subsystem is configured correctly. Other possible states could be:

NONE

The controller is not defined or is offline.

RESET

The controller is in the reset state.

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac1-hdisk2
dac0-hdisk3
dac1-hdisk4
dac0-hdisk5
dac1-hdisk6
dac0-hdisk7
dac1-hdisk8
```

Using the lsattr command

You can also use the **lsattr** command to find out which controller owns a particular hdisk. The following example shows portions of the outputs from several **lsattr** commands.

```
# lsattr -El dac0
passive_control no                Passive controller      False
alt_held_reset no                Alternate held in reset False
controller_SN 1T01710113           Controller serial number False
ctrl_type 3552                   Controller Type         False
cache_size 512                   Cache Size in MBytes   False
scsi_id 0x11100                SCSI ID                 False
lun_id 0x0                       Logical Unit Number    False

# lsattr -El dac1
passive_control no                Passive controller      False
alt_held_reset no                Alternate held in reset False
controller_SN 1T03910039           Controller serial number False
ctrl_type 3552                   Controller Type         False
cache_size 512                   Cache Size in MBytes   False
scsi_id 0x11000                SCSI ID                 False
lun_id 0x0                       Logical Unit Number    False

# lsattr -El hdisk1
----- Parts removed:
scsi_id 0x11100                SCSI ID                 False
lun_id 0x0000000000000000         Logical Unit Number    False
cache_method fast_write           Write Caching method   True
```

In the example, hdisk1 belongs to the controller represented by dac0. The hdisk has the same SCSI ID as the controller (dac) that owns it.


Identifying logical drives by operating system device names

The **lsattr** command provides detailed information about a volume, including information that allows you to map the system device name to the logical volume on the FASiT storage subsystem.

For example, if you run the **lsattr** command on the LUN named hdisk4, it provides the following information: it is a 6 GB LUN of type RAID, with a LUN ID of 2, and an IEEE volume name of 600A0B80000C1E6300000033BB39927 (see the following example). You can make a quick identification by locating the LUN ID on the far right side of the **Mappings View** tab.

```
# lsattr -El hdisk4
pvid           none                Physical volume identifier      False
q_type         simple              Queuing Type                   False
queue_depth    30                 Queue Depth                    True
reserve_lock   yes                 RESERVE device on open        True
write_cache    yes                 Write Caching enabled          True
size           6062               Size in Mbytes                 False
raid_level     0                  RAID Level                     False
rw_timeout     30                 Read/Write Timeout value      True
reassign_to    120                Reassign Timeout value        True
scsi_id        0x11000            SCSI ID                         False
lun_id         0x000200000000000 Logical Unit Number            False
cache_method   fast_write          Write Caching method          True
prefetch_mult  0                  Multiple of blocks to prefetch True
ieee_volname   600A0B80000C1E63000000033BB39927 IEEE Unique volume name      False
```

You can make a more exact correlation using the distinctive *ieee_volname* attribute. The value of this attribute on the AIX host is the same as the Unique Logical Drive Identifier on the FAS*T* storage subsystem. The **ieee_volname** and Unique Logical Drive Identifier can be found in the Change Mappings window, which is shown in Figure 6. (Click the **Mappings** view tab. Select **Logical Drive Name**. Right click and select **Change Mappings**.)



SJ000712

Figure 6. Change Mappings window

Another way to identify the system disk name based on the logical volume name is by using the Logical Drive Property window. Open this window by clicking the Storage Management window and right-clicking **Logical Drive Name** → **Properties**. Look for the **World-wide name**, **Capacity**, and **RAID level** properties.

The following example shows a portion of the output of the **lsattr** command for **hdisk20**. The *ieee_volname* and *lun_id* attributes have the same values as the **World-wide name** and **Subsystem ID (SSID)** properties that are displayed in the Logical Drive Property window, respectively.

```
# lsattr -El hdisk20
write_cache  yes                Write Caching enabled          True
size         9216                Size in Mbytes                 False
raid_level   5                    RAID Level                     False
scsi_id      0x210513                   SCSI ID                        False
lun_id       0x0012000000000000    Logical Unit Number           False
cache_method fast_write                   Write Caching method          True
prefetch_mult 0                    Multiple of blocks to prefetch on read True
ieee_volname 600A0B80000C1E63000000163BB39B7C IEEE Unique volume name       False
```

Identifying FlashCopy volume types

Perform the following steps to identify FlashCopy volume types:

1. Run the following command to list available volumes:

```
# lsdev -Cc disk
```

A list of available volumes appears, as shown in the following example:

```
[root@el1m17c171] / # lsdev -Cc disk |pg
hdisk0 Available 40-60-00-4,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 40-60-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 11-08-01 3552 (500) Disk Array Device
hdisk3 Available 11-08-01 3552 (500) Disk Array Device
hdisk4 Available 11-08-01 3552 (500) Disk Array Snapshot
```

The FlashCopy array is identified by the term **Disk Array Snapshot**.

2. Run the following command to display the state of each controller in a FASTt array:

```
# fget_config -Av
```

A list appears, as shown in the following example:

```
[root@el1m17c171] / # fget_config -Av
---dar0---
User array name = 'FASTt System 1'
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk2 Raid-0-0B
dac1-hdisk3 Raid-0-1B
dac1-hdisk4 Raid-3-0A-1 Snapshot Volume
```

The FlashCopy volume is identified by the term **Snapshot Volume**.

For more information about the **lsdev** and **fget_config** commands, see “Identifying device names and bus numbers” on page 30.

Using dynamic capacity expansion and dynamic volume expansion with AIX

Dynamic volume expansion (DVE) is dynamic on the FAS*T*, but it requires manual intervention for AIX to recognize the new volume capacity. This section explains how to use DVE with AIX.

Before you begin: Ensure that there is available free capacity within the array. You can check free capacity availability using the SMclient, in the Logical/Physical view. If there is not enough free capacity, and extra drives are available, perform a dynamic capacity expansion (DCE) operation before you perform the DVE operation. A DCE operation increases the capacity of the array by adding physical disks.

Performing a dynamic capacity expansion operation

Complete the following steps to perform a dynamic capacity expansion:

1. From the Logical/Physical view of the SMclient, right click on **Array** → **Add Free Capacity (Drives)**.

Note: The process might take several hours to complete. This process must be complete before you can proceed to the next step.

2. From the Logical/Physical view of the SMclient, right click on **Logical Drive** → **Increase Capacity**.
3. Type the amount that you want to increase the logical volume.

You will see an hourglass on every logical drive within the array. You must wait for the process to complete before you can begin any AIX intervention.

Note: If the FAS*T* is very busy, the process might take several hours to complete.

After you have ensured that there is sufficient free capacity within the array, you can perform DVE.

Note: You must run either the **reboot** or the **cfgmgr** command before you can view the additional capacity using the ODM attribute.

Performing a dynamic volume expansion (DVE) operation

If the host is running on AIX 5.2, use the following command to perform a dynamic volume expansion:

```
chvg -g volume group name
```

where the *volume group name* is the name of the associated LVM volume group.

If the host is running on AIX 4.3.3 or AIX 5.1, complete the following steps to perform a dynamic volume expansion:

1. Stop the associated LVM volume group by typing the following command:

```
varyoffvg
```

2. Put the associated hdisk into the defined state by typing the following command:

```
rmdev -1
```

3. Make the hdisk available again by typing the following command:

```
cfgmgr
```

You can view the increased capacity by typing the following command:

```
lsattr -El hdiskXX
```

4. Vary on the volume by typing the following command:

```
varyonvg
```

A message informs you of a capacity increase.

5. Print a list of the available physical partitions within the LVM volume group by typing the following command:

```
lsvg -p volume group name
```

where the *volume group name* is the name of the associated LVM volume group.

Note: If you see a warning that the volume group cannot be imported to AIX 5.1 or lower, ignore the warning. It is an error.

Figure 7 shows a sample output from this procedure.

```
# lsvg -p Blast50
Blast50:
PV_NAME    PV STATE    TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk50    active      127         76         25..00..00..25..26

# umount /blast/tfs50
# varyoffvg Blast50
# rmdev -l hdisk50
hdisk50 Defined
# cfgmgr -v
# varyonvg Blast50
0516-1434 varyonvg: Following physical volumes appear to be grown in size.
          Run chvg command to activate the new space.
          hdisk50
# chvg -g Blast50
0516-1224 chvg: WARNING, once this operation is completed, volume group Blast50
          cannot be imported into AIX 510 or lower versions. Continue (y/n) ?
y
# lsvg -p Blast50
Blast50:
PV_NAME    PV STATE    TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk50    active      255         204        25..26..51..51..51
```

Figure 7. Volume group Blast50 is increased from 4 GB to 8 GB

Resolving disk array errors

This section shows a list of possible disk array errors that could be reported in the AIX error log. You can view the AIX error log by running the **errpt -a** command.

You can also check your Storage Manager Major Event log (MEL) to find out whether there is any correlation between the host, SAN, and FAStT storage subsystem.

You might need to validate your configuration or replace defective hardware to correct the situation.

- **FCP_ARRAY_ERR1** *ARRAY OPERATION ERROR*
A permanent hardware error involving the disk array media.
- **FCP_ARRAY_ERR2** *ARRAY OPERATION ERROR*
A permanent hardware error.
- **FCP_ARRAY_ERR3** *ARRAY OPERATION ERROR*
A permanent error detected by the array adapter.
- **FCP_ARRAY_ERR4** *ARRAY OPERATION ERROR*
A temporary error within the array, communications, adapter, and so on.
- **FCP_ARRAY_ERR5** *UNDETERMINED ERROR*
- **FCP_ARRAY_ERR6** *SUBSYSTEM COMPONENT FAILURE*
A degradation condition has occurred other than a disk drive.
- **FCP_ARRAY_ERR7** *CONTROLLER HEALTH CHECK FAILURE*
A health check on the passive controller has failed.
- **FCP_ARRAY_ERR8** *ARRAY CONTROLLER SWITCH*
- **FCP_ARRAY_ERR9** *ARRAY CONTROLLER SWITCH FAILURE*
An array controller switch has failed.
- **FCP_ARRAY_ERR10** *ARRAY CONFIGURATION CHANGED*
A logical unit has been moved from one controller to the other (most likely by the action of an alternate host).
- **FCP_ARRAY_ERR11** *IMPROPER DRIVE TYPE FOR DUAL ACTIVE MODE*
This error should not be possible on the 2102 array, and exists for history reasons only. FCP_ARRAY_ERR11 may be reused for a different error in the future.
- **FCP_ARRAY_ERR12** *POLLED AEN FAILURE*
An automatic error notification has failed.
- **FCP_ARRAY_ERR13** *ARRAY INTER-CONTROLLER COMMUNICATION FAILURE*
- **FCP_ARRAY_ERR14** *ARRAY DRIVE FAILURE*
- **FCP_ARRAY_ERR15** *CACHE BATTERY LOW/DATA LOSS POSSIBLE*
If a controller card is replaced, it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the error log. Do not replace the controller.
- **FCP_ARRAY_ERR16** *CACHE BATTERY CHARGE BELOW 87.5%*
If a controller card is replaced, it is likely that the cache batteries will be flat. It can take two days for the cache batteries to be fully recharged. During this time errors are logged in the errorlog. Do not replace the controller.
- **FCP_ARRAY_ERR17** *WORLDWIDE NAME CHANGED*
A controller has changed worldwide names (most likely either it was replaced without placing it in the reset state first, or the cabling was changed so that a different controller with the same SCSI ID is on the loop).
- **FCP_ARRAY_ERR18** *RESERVATION CONFLICT*

An operation failed because the disk array volume (LUN) is reserved by another host.

Redistributing volumes in case of failure

If a failure occurs that initiates a controller failover, perform the following steps to redistribute logical drives to their preferred paths:

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage FAStT Installation and Support Guide* for the appropriate FAStT fibre-channel storage server.
2. Redistribute volumes to their preferred paths by clicking **Subsystem Management** → **Storage Subsystem** → **Redistribute Logical Drive**.
3. Run the **fget_config** command to verify the active paths, as shown in this example:

```
# fget_config -l dar0
dac0 ACTIVE dac1 ACTIVE
dac0-hdisk1
dac0-hdisk2
dac0-hdisk3
dac1-hdisk4
dac1-hdisk5
dac1-hdisk6
dac1-hdisk7
dac0-hdisk8
```


Chapter 3. Installing storage management station software on HP-UX systems

This chapter provides the following specific information for HP-UX operating systems:

- Hardware and firmware requirements
- Installing client software
- Installing host software
- Performing the initial configuration
- Creating direct-attached and SAN-attached configurations
- Enabling multipath I/O with PV-links
- Identifying the controller numbers
- Identifying device names and bus numbers

FAST600:

HP-UX operating systems are not supported on the FAST600 Fibre Channel Storage Server.

Hardware and firmware requirements

For system-level hardware and software requirements, see Appendix B, “HP-UX system requirements”, on page 141.

Table 14 lists the supported versions of hardware to use with Storage Manager 8.3.

Table 14. Supported versions of hardware for HP-UX systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAST900 RAID Controller Enclosure Unit	1742-90U	Appware 05.30.07.00 NVSRAM N1742F900R830V03
IBM TotalStorage FAST700 RAID Controller Enclosure Unit	1742-1RU	Appware 05.30.07.00 NVSRAM N1742F700R830V03
IBM TotalStorage FAST700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9321
IBM FAST500 RAID Controller Enclosure Unit	3552-1RU	Appware 05.30.07.00 NVSRAM N3552F500R830V03
IBM FAST500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9166
IBM FAST200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05
IBM FAST200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05
HP FC Adapter	A5158A	B.11.00.06 (driver)
HP FC Adapter 2 Gb	A6795A	B.11.00.01 (driver)
Brocade switch	2109-S08 2109-S16	2.6.0j
Brocade switch	2109-F16	3.0.2h

Table 14. Supported versions of hardware for HP-UX systems (continued)

Product Name	Model	Product release and firmware version
Brocade switch	2109-M12, 2109-F32	4.0.2b
Brocade switch	2109-F08	3.0.2k
McData switch	2032-064, 2032-140	4.1
McData switch	2031-016, 2031-216, 2031-224	4.1
InRange switch	2042-001	4.1.05
Cisco switch	2062-D01, 2062-D07, 2062-T07	GA level

Note: Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. Read the readme.txt file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

Installing the client software

Use the procedures in this section to install the client software on an HP-UX storage management station or on an HP-UX host acting as a storage management station. Install the software in the following order:

1. SMruntime
2. SMclient

Note: SMclient is dependent on SMruntime, which is a Java compiler for the SMclient and must be installed first. If you do not install SMruntime you will receive errors during the software install analysis phase.

Prerequisites

Before installing the client software, ensure that the following conditions are met:

- This is the HP 9000-series server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in Appendix B, "HP-UX system requirements", on page 141.
- Neither SMruntime nor SMclient is installed on the host and you are ready to install them on this machine.

Installing SMruntime

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the SMruntime installation process.

1. Insert the HP-UX installation CD in the CD-ROM drive.
2. Start the installation process by typing the following command:

```
# mount -o cdfcase -F cdfs /device_pathname /cdrom
# swinstall -s /cdrom/HP-UX/HPruntime_0804500
```

The Software Selection window opens and displays this message:

```
SMruntime
```

3. Select **SMruntime**.
4. Select **Actions** → **Mark for installation**.
5. Select **Actions** → **Install (Analysis)**.
6. Click **OK** when the analysis is finished.
7. In the confirmation window, click **Yes** to start the installation.
8. When the application is finished, click **Done** in the Installation window and close the application.

Note: You must exit the **swinstall** command before running the **swverify** command to verify the installation.

9. Verify that the installation was successful by typing the following command:

```
# swverify -v SMruntime
```

If no failure is reported, go to step 10. If a failure is reported, follow the instructions in the /var/adm/sw/swagent.log file.

10. You have completed the SMruntime installation of this HP-UX storage management station.

Installing the SMclient

In the following procedure, the installation CD is mounted at /cdrom. Adjust these instructions as required for your specific installation. No restart is required during the SMclient installation process.

1. Start the installation process by typing the following command:

```
# swinstall -s /cdrom/HP-UX/SMclient-HP-081G500
```

The Software Selection window opens and displays the following message:

```
SMclient
```

2. Select **SMclient**.
3. Select **Actions** → **Mark for installation**.
4. Select **Actions** → **Install (Analysis)**.
5. Click **OK** when the analysis is finished.
6. In the Confirmation window, click **Yes** to start the installation.

7. When the application is finished, click **Done** in the Installation window and close the application.

Note: You must exit the **swinstall** command before running the **swverify** command to verify installation.

8. Verify that the installation was successful by typing the following command:

```
# swverify -v SMclient
```

If no failure is reported, go to step 9. If a failure is reported, follow the instructions in the `/var/adm/sw/swutil.log` file.

9. You have completed the SMclient installation of this HP-UX storage management station.

Note: To ensure redundancy in a cluster environment, you must install SMclient on at least one additional storage management station or cluster server.

Installing the host software

Use the procedures in this section to install the host software on an HP-UX storage management station or on an HP-UX host acting as a storage management station. Install the software in the following order:

1. SMagent (optional — for in-band management only)
2. SMutil

Note: SMagent and SMutil are dependent on SMruntime, which is a Java compiler for the storage manager software and must be installed first. If you do not install SMruntime you will receive errors during the software install analysis phase.

Prerequisites

Before installing the host software, ensure that the following conditions are met:

- This is the HP 9000-series server that you identified as the storage management station.
- This machine meets the minimum hardware and software requirements described in Appendix B, “HP-UX system requirements”, on page 141.
- SMruntime has been installed.

Installing the SMagent (optional)

SMagent is required for in-band management only. It is not required for out-of-band management.

In the following procedure, the installation CD is mounted at `/cdrom`. Adjust these instructions as required for your specific installation. No restart is required during the SMagent installation process.

1. Start the installation process by typing the following command:

```
# swinstall -s /cdrom/HP-UX/SMagent-HP-08204500
```

The Software Selection window opens and displays this message:

```
SMagent
```

2. Select **SMagent**.
3. From the toolbar, click **Actions** → **Mark for installation**.
4. Click **Actions** → **Install (analysis)**.
5. Click **OK** when the analysis is finished.
6. In the Confirmation window, click **Yes** to start the installation.
7. When the installation is finished, click **Done** in the Installation window and close the application.

Note: You must exit the **swinstall** command before running the **swverify** command to verify the installation.

8. Verify that the installation was successful by typing the following command:

```
# swverify -v SMagent
```

If no failure is reported, go to step 9. If a failure is reported, follow the instructions in the `/var/adm/sw/swagent.log` file.

9. You have completed SMagent installation on this HP-UX storage management station.

Installing SMutil

In the following procedure, the installation CD is mounted at `/cdrom`. Adjust these instructions as required for your specific installation. No restart is required during the SMutil installation process.

1. Start the installation process by typing the following command:

```
# swinstall -s /cdrom/HP-UX/SMutil_08204500
```

The Software Selection window opens and displays this message:

```
SMutil
```

2. Select **SMutils**.
3. Select **Actions** → **Mark for Installation**.
4. Select **Actions** → **Install (Analysis)**.
5. Click **OK** when the analysis is finished.
6. In the Confirmation window, click **Yes** to start the installation.
7. When the application is finished, click **Done** in the Installation window and close the application.

Note: You must exit the **swinstall** command before running the **swverify** command to verify the installation.

8. Verify that the installation was successful by typing the following command:

```
# swverify -v SMutil
```

If no failure is reported, go to step 9 on page 46. If a failure is reported, follow the instructions in the `/var/adm/sw/swutil.log` file.

9. You have completed the SMutil installation of this HP-UX storage management station.

Performing the initial configuration of storage subsystems on HP-UX hosts

After you install the Storage Manager 8.3 software you need to configure the subsystem. System configuration includes configuring logical drives and creating storage arrays through the Subsystem Management window of the SMclient application.

See “Configuring storage subsystems” on page 74 for additional information about storage configuration.

Before you begin: Note the following information:

- All FAStT storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31.
- AVT/ADT is enabled by default; leave it in that state for HP-UX hosts.
- Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. Read the readme.txt file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

- If you have purchased a Storage Partitioning Premium Feature, ensure that the feature is enabled. See “Enabling Storage Partitioning” on page 10 for more information.

Complete the following procedures to configure Storage Manager 8.3 for an HP-UX system.

Steps for adding storage subsystems to SMclient

1. To set up the storage subsystem for HP-UX, the subsystem must be physically configured, at least initially, for direct management through the Ethernet connections on each controller as well as through the fibre-channel connection. Install the Storage Manager 8.3 software (SMruntime, SMclient, SMagent, and SMutil) before configuring the subsystem.
2. After the disk subsystem is configured on the network, type the following command to run SMclient on either the host server, if it is on the same network as the storage subsystem, or on another machine in the network that has the Storage Manager 8.3 software installed:

```
# SMclient
```

3. Complete the following steps to specify the IP addresses of the controllers:
 - a. In the Enterprise Management window, click **Edit** → **Add Device**.
 - b. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.
 - c. Type the IP address of the second controller and click **Add**, and then click **Close**.

Note: Failure to add both controllers results in a partially-managed system.

4. In the Subsystem Management window, click **Storage Subsystem** → **Profile** and review the summary portion of the output.

Steps for updating NVSRAM and FASTT firmware

Complete the following steps to upgrade the NVSRAM:

1. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.
2. Click **Storage Subsystem** → **Download** → **NVSRAM**.
3. In the NVSRAM window, go to the directory where the latest NVSRAM file resides.
4. Type or select the full pathname of the NVSRAM file, and then click **Update** → **OK**.

If required, the host restart process starts after the storage subsystem has been turned off.

Complete the following steps to upgrade the firmware:

1. Click **Storage Subsystem** → **Download** → **Firmware**.
2. Go to the directory where the firmware resides.
3. Type or select the full pathname of the firmware file, and then click **Update** → **OK**.

The firmware is downloaded to the controllers. When the download is finished, the storage subsystem becomes unresponsive and you are returned to the Enterprise Management window.

Type the following command to restart the host:

```
# reboot
```

Setting up an HP-UX host group

Note: Before you set up an HP-UX host group, identify the worldwide port name for each of your HBAs. Look for the worldwide port name for the HBA and write down the number for future reference. Use the following command to find your HBAs listed in the dev directory as td0, td1, and so on, where X in the following example represents the td number.

```
#cd /dev/td  
#tdutil /dev/tdX
```

1. Click the **Mappings View** tab on the Subsystem Management window.
2. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.
3. Type the name of the new host group (for example, HP-UX). Click **Add**, and then click **Close**.
4. Select the new host group and click **Mappings** → **Define** → **Host**.
5. Define the new host. Type the name of the HP-UX host to which the storage subsystem is attached.
 - a. Click **Add**, and then click **Close**.

- b. Select the host that you just added and right-click **Define New Host Port**.

Note: You might be required to restart the host if you cannot see the WWPN for your HBAs. Restarting causes the system to log into the switch.

- c. Select the desired host port for the first HBA, and then change the host type to **HP-UX** and click **Add**.

Note: Failure to change the host type from the default to HP-UX causes undesired results. Also, if you use two HBAs you must define two host ports. Failure to do so causes undesired results.

- d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

Creating direct-attached and SAN-attached configurations

Storage Manager 8.3 supports FASSt200, FASSt500, FASSt700, and FASSt900 storage servers in direct-attached HP-UX configurations or in a SAN environment through switches in HP-UX configurations.

Creating a direct-attached configuration

To create a direct-attached configuration, you must ensure that:

- One or two HP-UX servers can be connected to the FASSt storage server:
 - FASSt200 can support one HP-UX server.
 - FASSt500 can support two HP-UX servers.
 - FASSt700 can support two HP-UX servers.
 - FASSt900 can support two HP-UX servers.
- Two server FASSt500, FASSt700 or FASSt900 configurations require four host-side mini-hubs, each with exactly one fibre-channel connection from each HBA to a mini-hub.
- There are two or four HBAs (HP A6795A or HP A5158A) per FASSt storage server. Each pair must be configured to one FASSt partition.
- No external hubs are being used.

To set up a direct-attached configuration, follow these steps:

1. Connect the HBAs to each controller or mini-hub port of the FASSt storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

To create a SAN-attached configuration, you must ensure that:

- Multiple fibre-channel HBAs within the same server cannot “see” the same FASSt controller port.
- The IBM fibre-channel HBAs are isolated from each other if they are connected to the same switch that is connected to the same FASSt controller port.
- Each fibre-channel HBA and controller port must be in its own fabric zone, if they are connecting through a single fibre-channel switch, such as a 2109-F16.

See the documentation that is provided by the switch manufacturer for more information about zoning. Multiple FASSt devices can be configured to the same set of fibre-channel HBAs through a fibre-channel switch.

To set up a SAN-attached configuration, follow these steps:

1. Connect the HBAs to the switch or switches.
2. Connect the FAStT storage subsystems to the switch or switches.
3. Set the required zones on the fibre-channel switch or switches, if applicable.

Note: For information about zoning and enabling zones, see the documentation that is provided by the switch manufacturer.

4. Configure and verify the configuration.

Enabling multipath I/O with PV-links

If the HP-UX system is attached with two host bus adapters to the FAStT storage server, you can establish redundant access to storage by using physical volume links (PV-links), a feature of the HP-UX operating system. PV-links achieve access redundancy by using devices with both primary and secondary paths to the same device.

Before you begin: Ensure that SMutil is installed on the host so that you can run the **hot_add** and **SMdevices** commands.

Use the following procedure to establish redundant access to storage by using PV-links:

1. Run the **hot_add** command from HP-UX at the shell prompt. This command updates any new devices that are created or added. A dump is generated. When the **hot_add** command runs, each new logical drive that is created in the Subsystem Management window represents a disk device to the operating system.

```
#hot_add
```

2. Run the **SMdevices** command. The system provides a dump similar to the example in Figure 8 on page 50. Notice that every logical drive or device and access logical unit number (LUN) is listed twice, because you have two paths to each logical drive.

```
#SMdevices
```

```

IBM FASTT Storage Manager Devices, Version 08.3.45.00
Built Wed Mar 20 00:58:59 GMT+00:00 2002
(C) Copyright International Business Machines Corporation, 2002
Licensed Material - Program Property of IBM. All rights reserved.

/dev/rdisk/c28t0d2 [Storage Subsystem 700var201, Logical Drive hp_1, LUN 1,
Logical Drive WWN <600a0b8000075a54000000803d2c6de6>, Preferred Path
(Controller-A): In Use]

/dev/rdisk/c28t0d3 [Storage Subsystem 700var201, Logical Drive hp_2, LUN 2,
Logical Drive WWN <600a0b8000075a540000006d3d2c6d9f>, Alternate Path
(Controller-A): In Use]

/dev/rdisk/c28t0d4 [Storage Subsystem 700var201, Logical Drive hp_3, LUN 3,
Logical Drive WWN <600a0b8000075a54000000813d2c6df4>, Preferred Path
(Controller-A): In Use]

/dev/rdisk/cc28t0d5 [Storage Subsystem 700var201, Logical Drive hp_4, LUN 4,
Logical Drive WWN <600a0b8000075a540000006e3d2c6da9>, Alternate Path
(Controller-A): In Use]

/dev/rdisk/c28t3d7 [Storage Subsystem 700var201, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b8000075a540000000200000000>]

/dev/rdisk/c30t0d2 [Storage Subsystem 700var201, Logical Drive hp_1, LUN 1,
Logical Drive WWN <600a0b8000075e60000000803d2c6de6>, Preferred Path
(Controller-A): In Use]

/dev/rdisk/c30t0d3 [Storage Subsystem 700var201, Logical Drive hp_2, LUN 2,
Logical Drive WWN <600a0b8000075e600000006d3d2c6d9f>, Alternate Path
(Controller-A): In Use]

/dev/rdisk/c30t0d4 [Storage Subsystem 700var201, Logical Drive hp_3, LUN 3,
Logical Drive WWN <600a0b8000075e60000000813d2c6df4>, Preferred Path
(Controller-A): In Use]

/dev/rdisk/c30t0d5 [Storage Subsystem 700var201, Logical Drive hp_4, LUN 4,
Logical Drive WWN <600a0b8000075e600000006e3d2c6da9>, Aternate Path
(Controller-A): In Use]

/dev/rdisk/c30t3d7 [Storage Subsystem 700var201, Logical Drive Access, LUN 31,
Logical Drive WWN
<600a0b8000075e600000000200000000>]

```

Figure 8. Device names

3. Determine the primary and alternate path for each logical drive. A part of the WWN of each logical drive is unique for each controller in the FASTT storage server. If you examine the WWNs for the access volumes in Figure 8, you will notice that they differ in only five digits, 75a54 and 75e60. The WWNs for the logical drives also differ in these five digits.

In the example in Figure 8, the devices are viewed through the adapters c28 and c30. To determine the primary path of a specific logical drive do the following:

- a. Find the WWN for each access LUN. In this case, Access LUN 1 is associated with c28 and has the WWN of 75a54. Access LUN 2 is associated with c30 and has the WWN of 75e60.

- b. Identify the primary path by matching the logical drive WWN to an access LUN WWN. In this case, the WWN for the LUN 10 (75a54) is associated with adapter c28. Therefore, the primary path would be c28t1d2 and the alternate path would be c30t1d2.
- c. For future reference, enter this path information into a primary and secondary path matrix, similar to the one in Table 15.

Table 15. Primary and secondary path matrix information

Logical drive	Primary path	Secondary path
hp_1	c28t1d2	c30t1d2

- d. Repeat these steps for each logical drive.
4. Define the primary paths by typing the following command. The system confirms the creation of the new physical volume.

```
#pvcreate /dev/rdisk/c28t0d0
```

5. Make a directory for volume group by typing the following command. You need the directory to reside in the /dev directory.

```
#cd /dev
#mkdir vg1
```

6. Make a node by typing the following command:

```
#mknod /dev/vg1/group c 64 0x01000
```

7. Create PV-links for the device with the primary path by typing the following command:

```
#vgcreate /dev/vg1 /dev/dsk/C28t1d2
```

8. Create file system with logical volumes.
9. Add an alternate path for the device by typing the following command:

```
#vgextend vg01 /dev/dsk/C30t1d2
```


10. Repeat steps 4 - 9 for each device.
11. Verify the primary and alternate paths for each device by typing the following command, where *vgname* is the volume group name.

```
#vgdisplay -v vgname
```

Identifying the controller ID numbers

Complete this procedure to identify the controller ID numbers:

1. Select a controller in the Physical View of the Subsystem Management window.
2. Click **Controller** → **Properties** → **Interfaces**. The Controller Properties window opens, as shown in Figure 9 on page 52.



SJ000702

Figure 9. Controller Properties window

The first number in the **Current Loop ID** field is the controller ID, which can be one or two digits. For hubs and directly-connected subsystems, the controller ID number and the switch port ID can be found in the same window by converting the middle hexadecimal number to a decimal number.

Identifying device names and bus numbers

Use the **ioscan** command to determine the full device names of data logical drives in the storage subsystem. This information allows you to determine the host adapter external bus numbers and the IDs of the controllers that the host adapters are connected to.

1. View the operating system list of devices by typing the following command:

```
# ioscan -f
```

2. The identification information for every device is displayed. See Figure 10 on page 53.

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
bc	0		root	CLAIMED	BUS_NEXUS	
bc	1	8	ccio	CLAIMED	BUS_NEXUS	I/O Adapter
fc	0	8/12	fcT1	CLAIMED	INTERFACE	HP Fibre Channel
Storage lan Storage Cntl				CLAIMED	INTERFACE	HP Fibre Channel
fc	0	8/12.8	fc		INTERFACE	FCP Protocol Adapter
ext_bus	2	8/12.8.0.2.0.0	fcpdev	CLAIMED	INTERFACE	FCP Device Interface
target	39	8/12.8.0.2.0.0.0	tgt	CLAIMED	DEVICE	
disk	130	8/12.8.0.2.0.0.0	sdis		DEVICE	IBM INF-01-00
disk	35	8/12.8.0.2.0.0.7	sdisk	CLAIMED	DEVICE	IBM Universal Xport
target	39	8/12.8.0.2.0.1	tgt	CLAIMED	DEVICE	
disk	131	8/12.8.0.2.0.1.0	sdisk	CLAIMED	DEVICE	IBM INF-01-00

SJ000701

Figure 10. Device identification information

- Examine the information. LUN numbers range from 0 - 7 and target addresses from 0 - 3.

Device names have the following format: c<x>t<y>d<z> where:

<x>

Represents the controller ID number

<y>

Represents the target address

<z>

Represents the LUN number

For example, disk 130 in Figure 10 has a device name of c2t0d0.

Redistributing volumes in case of failure

Auto Volume Transfer (AVT) is enabled, by default, on HP-UX hosts. If a failure occurs that initiates a controller failover, AVT redirects I/O to the available controller. AVT does not require manual redistribution.

Important: If you have a heterogeneous environment, you must redistribute logical volumes manually. Failure to do so will leave the subsystem in a *Needs Attention* state, because hosts that do not support AVT or have AVT disabled will not automatically redirect I/O to the preferred controller. If a failure occurs that initiates a controller failover, perform the following steps to redistribute logical drives to their paths:

- Repair or replace any faulty components. For more information, see the *IBM TotalStorage FASTT Installation and Support Guide* for the appropriate FASTT fibre-channel storage server.
- Redistribute volumes to their preferred paths by clicking **Subsystem Management** → **Storage Subsystem** → **Redistribute Logical Drive**.

Chapter 4. Installing storage management station software on Solaris systems

This chapter provides the following specific information for Solaris operating systems:

- Hardware and firmware requirements
- Installing client software
- Installing host software
- Performing the initial configuration
- Partitioning (default)
- Creating direct-attached and SAN-attached configurations
- Configuring JNI HBA cards
- Installing RDAC
- Verifying external storage

Hardware and firmware requirements

Table 16 lists the FAStT storage servers that are supported with Storage Manager 8.3.

Table 16. FAStT storage servers that are supported on Solaris systems

Product Name	Model	Product release and firmware version
IBM TotalStorage FAStT900 RAID Controller Enclosure Unit	1742-90U	Appware 05.30.09.00 NVSRAM N1742F900R830V04
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 05.30.09.00 NVSRAM N1742F700R830V03
IBM TotalStorage FAStT600 RAID Controller Enclosure Unit	1722-60U, 1722-60X	Snapware 08.33.05.00 NVSRAM N1722F600R833V02
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 05.30.09.00 NVSRAM N3552F500R830V03
IBM FAStT200 RAID and Storage Unit, Single Controller	3542-1RU and FC 2101	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.09.00 NVSRAM N3542-2RUR830V05

Table 17 lists the FAStT storage servers that are supported with Storage Manager 8.3.

FAStT600: The FAStT600 does not support all of the expansion units, switches and JNI cards that are listed in Table 17. Please note the checkmarks that indicate which versions are supported on the FAStT600.

Table 17. Supported versions of hardware for Solaris systems

Product Name	Model	Product release and firmware version	Supported on FAStT600
IBM TotalStorage FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9321	✓

Table 17. Supported versions of hardware for Solaris systems (continued)

Product Name	Model	Product release and firmware version	Supported on FAStT600
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9166	
JNI FC Adapter	FC64-1063	fcaw2.5.18.pkg	
JNI FC Adapter	FCE-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1	
JNI FC Adapter	FCE-1473 S-bus 64 bit	JNIC146x.pkg, version 5.2.1	✓
JNI FC Adapter	FCE-6410-N 33 Mhz	JNIC.pkg, version 4.1.1.1	
JNI FC Adapter	FCE-6460-PCI 64 bit	JNIC146x.pkg, version 5.2.1	✓
JNI FC Adapter	FCE2-1063 S-bus 64 bit	JNIC.pkg, version 4.1.1.1	
JNI FC Adapter	FCE2-6412 66 Mhz 64 bit	JNIC.pkg, version 4.1.1.1	
JNI FC Adapter	FCE-1063	fca-pci.2.5.18.pkg	
Brocade switch	2109-S08, 2109-S16	2.6.0j	
Brocade switch	2109-F16	3.0.2k	✓
Brocade switch	2109-F8	3.0.2k	✓
Brocade switch	2109-M12, 2109-F32	4.0.2b	✓
McData switch	2032-064, 2032-140	4.1	✓
McData switch	2031-016	4.1	
McData switch	2031-216, 2031-224	4.1	✓
InRange switch	2042-001	4.1.05	✓
Cisco switch	2062-D01, 2062-D07, 2062-T07	GA level	✓

Note: Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. Read the readme.txt that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware, NVSRAM, disk drive firmware, and host adapter device drivers:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

For system-level hardware and software requirements, see Appendix C, “Solaris system requirements”, on page 143.

Installing the client software

Use the procedures in this section to install the client software on a Solaris storage management station.

Install the software in the following order:

1. SMruntime
2. SMclient

Note: The SMclient is dependent on SMruntime, which is a Java compiler for the SMclient and must be installed first.

Prerequisites

Before installing the client software, ensure that the following conditions are met:

- This is the SUN SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 2.7, or 8 and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements”, on page 143.
- Neither SMruntime nor SMclient is installed on the host and you are ready to install them on this machine.

Installing SMruntime

In the following procedure, the installation CD is mounted at /cdrom/sm83. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

Before you begin: Visit the Web sites that are listed in “Hardware and firmware requirements” on page 55 to ensure you have the most recent version of the software.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Start the installation process by typing the following command:

```
pkgadd -d /cdrom/SM83/SUN_SParc/SM8runtime-Sparc-filename.pkg
```

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:  
1 SMruntime  
IBM FAStT Storage Manager 8 Runtime (sparc)  
version number  
Select package(s) you wish to process (or 'all' to process all  
Packages). (default:all) [?,??,q]:
```

3. Type the value of the package that you are installing and press Enter. The installation process begins. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user  
Permission during the process of installing this package.  
Do you want to continue with the installation of <SMruntime>  
[y, n, ?]
```

4. Type `y` and press Enter. The installation process continues. When the SMruntime software has been successfully installed, the following message is displayed:

```
Installation of <SMruntime> was successful.
```

5. Type the following command to verify that the installation was successful:

```
#pkgchk Smruntime
```

The installation was successful if no output is displayed on the screen.

Installing SMclient

Use the following procedure to install and integrate the client software on a Solaris storage management station.

Note: SMclient is supported on Solaris 2.6, 2.7, and 8.

In the following procedure, the installation CD is mounted at `/cdrom/sm83`. Adjust these instructions as required for your specific installation location. No restart is required during the client installation process.

1. Insert the Solaris installation CD in the CD-ROM drive.
2. Start the installation process by typing the following command:

```
# pkgadd -d /cdrom/SM83/SUN_Sparc/SM8client-Sparc-filename.pkg
```

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might see displayed.

```
The following packages are available:
```

```
1 SMclient                IBM FAStT Storage Manager 8 Client  
                        (sparc) version number
```

```
Select package(s) you wish to process (or 'all' to process all  
packages). (default:all) [?,??,q]:
```

3. Type the value of the package that you are installing and press Enter.
The installation process begins. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user  
permission during the process of installing this package.
```

```
Do you want to continue with the installation of <SMclient>
```

```
[y, n, ?]
```

4. Type `y` and press Enter.

The installation process continues. When the SMclient software has been successfully installed, the following message is displayed:

```
Installation of <SMclient> was successful.
```

5. Type the following command to verify that the installation was successful:

```
# pkgchk SMclient
```

The installation was successful if no output is displayed on the screen.

6. You are finished with the client software installation on this Solaris storage management station. Remove the installation CD from the CD-ROM drive.

Note: To ensure redundancy in a cluster environment, you must install the client software on at least one additional storage management station or cluster server. To install the client software additional storage management stations, repeat step 1 on page 58 through step 6 of the installation procedure. Use the corresponding installation profile for each storage management station as a guide.

Installing host software on Solaris hosts

After you install the client software and configure your storage subsystems, use these instructions to install the appropriate host software.

The host software for Solaris consists of the following packages:

- SMagent (optional - for in-band management only)
- SMutil
- RDAC

This section contains the procedure for installing SMagent and SMutil. For information about how to install RDAC, see “Installing RDAC on Solaris hosts” on page 69.

Prerequisites

Before installing the host software, ensure that the following conditions are met:

- This is the Sun SPARCstation workstation that you have identified as the storage management station.
- This machine is running Solaris 2.6, 7, or 8, and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements”, on page 143.
- The RDAC, SMagent, and SMutil software packages are not installed and you are ready to install them on this machine.

Installing SMagent (optional)

SMagent is required for in-band management only. It is not required for out-of-band management.

Perform the following procedure to install SMagent:

1. Type the following command to start installing the SMagent package:

```
pkgadd -d /cdrom/SM83/SUN_SParc/SM8agent-Sparc-filename.pkg
```

The installation process begins.

Information about packages that can be installed in the specified directory is displayed on the command line. The following is an example of what you might

see displayed.

```
The following packages are available:

1 SMagent                IBM FAStT Storage Manager 8 Agent
                        (sparc) version number

Select package(s) you wish to process (or 'all' to process all
packages). (default:all) [?,??,q]:
```

2. Type the value of the package that you are installing and press Enter. The installation process begins. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user
Permission during the process of installing this package.
Do you want to continue with the installation of <SMagent>
[y n, ?]
```

3. Type y and press Enter. The installation process continues. When the SMagent software has been successfully installed, the following message is displayed:

```
Installation of <SMagent> was successful.
```

4. Type the following command to verify that the installation was successful:

```
# pkgchk SMagent
```

The installation was successful if no output is displayed on the screen.

5. Remove the installation CD from the CD-ROM drive.

Installing SMutil

Perform the following steps to install SMutil:

1. Type the following command to start installing the SMutil package:

```
pkgadd -d ./pathname/SMutil-Sparc-filename.pkg
```

The installation process begins.

Information about packages that you can install in the specified directory is displayed on the command line. An example of what you might see follows.

```
The following packages are available:

SMutil IBM FAStT Storage Manager 8 Util
                        (sparc)version number

Select package(s) you wish to process (or 'all'
to process all packages). (default:all) [?, ??, q]:
```

2. Type the value of the package that you are installing and press Enter. The installation process begins. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user
Permission during the process of installing this package.
Do you want to continue with the installation of <smutil>
[y, n, ?]
```

3. Type `y` and press Enter.
4. Type the following command to verify that the installation was successful:

```
# pkgchk SMagent
```

The installation was successful if no output is displayed on the screen.

When the SMutil has been successfully installed, the following message is displayed:

```
Installation of <smutil> was successful.
```

You have finished installing the host software on this Solaris host.

Uninstalling Storage Manager 8.3 software

When you uninstall SMruntime, RDAC, SMclient, SMagent, or SMutil software, the directory `/var/opt/SM8` might not be removed. Delete this directory to ensure a clean installation of the new software.

Performing the initial configuration of storage subsystems on Solaris hosts

To configure Storage Manager 8.3 for a Solaris system, complete the following procedures in the order that they are described in this section:

- Steps for adding storage subsystems to SMclient
- Steps for updating FAStT firmware and NVSRAM
- Steps for setting up a Solaris host group

Before you begin: Read the following information:

- All FAStT storage servers ship with NVSRAM configured for Windows NT hosts and have an access volume set to LUN 31. This should not present any problems for Solaris in recognizing the storage subsystem.
- If you have purchased a Storage Partitioning Premium Feature, ensure that the feature is enabled. See “Enabling Storage Partitioning” on page 10 for more information.
- Versions of firmware and NVSRAM that ship with this product might be later releases than those described in this document. Read the `readme.txt` file that is shipped with the product and go to one of the following Web sites to ensure that you have the latest versions of the firmware and the NVSRAM:

```
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt200
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900
```

Note: If you do not have the latest versions of firmware and NVSRAM, download them and perform the procedure described in “Steps for updating NVSRAM and FASTT firmware”. If the version numbers are current, you can skip that procedure.

Steps for adding storage subsystems to SMclient

Complete the following steps to specify the IP addresses of the controllers:

1. In the Enterprise Management window, click **Edit** → **Add Device**.
2. In the Add Device window, type the IP address of the first controller in the storage subsystem and click **Add**.
3. Type the IP address of the second controller and click **Add**, and then click **Close**.

Steps for updating NVSRAM and FASTT firmware

Complete the following steps to upgrade the NVSRAM and then upgrade the firmware:

1. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **NVSRAM**.
2. In the NVSRAM window, go to the directory where the latest NVSRAM file resides.
3. Type or select the full pathname of the NVSRAM file, and then click **Update** → **OK**.
4. Click **Storage Subsystem** → **Download** → **Firmware**.
5. Type or select the full pathname of the firmware file, and then click **Update** → **OK**.

The firmware is downloaded to the controllers. When the download is finished, the firmware becomes unresponsive and you are returned to the Enterprise Management window.

Steps for setting up a Solaris host group

Complete the following steps to set up a Solaris host group.

Note: In a cluster partition, perform logical drive mappings on the host group level so that all the hosts can see the same storage. In a normal partition, perform logical drive mappings on the host level.

1. Rescan the storage subsystem, and then click **Manage the Device** to return to the Subsystem Management window.
2. In the Subsystem Management window, click the **Mappings View** tab.
3. In the Mappings window, create a new host group by clicking **Mappings** → **Define** → **Host Group**.
4. Type the name of the new host group (for example, Solaris). Click **Add**, and then click **Close**.
5. Select the new host group and click **Mappings** → **Define** → **Host**.
6. Define the new host. Type the name of the Solaris host to which the storage subsystem is attached.
 - a. Click **Add**, and then click **Close**.
 - b. Right-click the host that you just added and select **Define New Host Port**.
 - c. Select the desired host port for the first HBA, and then change the host type to **Solaris** and click **Add**.

Note: Failure to change the host type from the default to Solaris causes undesired results.

- d. Choose the host port for the second HBA and click **Add**, and then click **Close**.

Note: If you use two HBAs you must define two host ports. Failure to do so causes undesired results.

- e. To view the added LUNs, click the host group, and then click **Mappings** → **Define** → **Define Additional Mapping**.
- f. Type the following command to run the hot_add utility:

```
# /etc/raid/bin/hot_add
```

For information about the output of the utility, see “Default partitioning for Solaris devices”.

- g. Click **File** → **Exit**.
- h. Obtain the device name for each volume (or LUN) by typing the following command:

```
# format
```

The SMclient is now able to connect to the storage subsystem.

Default partitioning for Solaris devices

After you install the Storage Manager 8.3 software and configure the subsystems and logical drives through the Subsystem Management window of the SMclient application, run the `/etc/raid/bin/hot_add` utility from a shell prompt. (For additional information about the `hot_add` utility see the applicable man page.)

When the `hot_add` utility is run, each logical drive that is created in the Subsystem Manager presents a disk device to the operating system. The devices have the following naming convention:

```
/dev/[r]dsk/c<X>t<X>d<X>s<X>
      |         |         |         |
      |         |         |         s=slice
      |         |         |         d=drive
      |         |         t=target
      |         c=controller
      disk subdirectory
```

where X represents a number. The number following the slice (s) represents the partition on the disk and can be from 0 - 7.

Note: To run the `hot_add` or `SMdevices` commands, you must have `SMutil` installed on the host.

You can overwrite disk information by using partitions or slices 0 - 2; this results in an inability to mount a disk. For this reason, do not use slices 0, 1, or 2. Using these slices can cause unpredictable results. Mount file systems only on slices 3, 4, 5, 6, or 7 on the devices that are derived from the logical drives on the storage subsystem.

Direct-attached and SAN-attached configurations

Storage Manager 8.3 supports FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 storage servers in direct-attached Solaris configurations or in a SAN environment through switches in Solaris configurations.

FAStT600:

For the FAStT600, you must set the link speed at 2 Gb on the controllers and on the FAStT700 EXP.

Creating a direct-attached configuration

In a direct-attached configuration, one or two Solaris servers can be connected to the FAStT storage server, as follows:

- FAStT200 can support one Solaris server.
- FAStT500, FAStT600, FAStT700, and FAStT900 can support two AIX servers.

Requirements:

- Two-server FAStT500, FAStT700, or FAStT900 configurations require four host-side mini-hubs, each with exactly one fibre-channel connection from each HBA to a mini-hub.
- There must be two or four JNI HBAs per FAStT storage server. Each pair must be configured to one FAStT partition.
- No external hubs can be used.

Perform the following steps to set up a direct-attached configuration:

1. Connect the HBAs to each controller or mini-hub port of the FAStT storage server.
2. Configure and verify the configuration.

Creating a SAN-attached configuration

Use the following procedure to create a SAN-attached configuration.

Requirements:

- Multiple HBAs within the same server must be unable to “see” the same FAStT controller port.
- The JNI HBAs must be isolated from each other if they are connected to the same switch that is connected to the same FAStT controller port.
- Each HBA and controller port must be in its own fabric zone, if they are connecting through a single switch, such as a 2109-F16.

See the documentation provided by the switch manufacturer for more information about zoning. Multiple FAStT devices can be configured to the same set of HBAs through a fibre-channel switch.

Perform the following steps to set up a SAN-attached configuration:

1. Connect the HBAs to the switch or switches.
2. Connect the FAStT storage subsystems to the switch or switches.
3. Set the required zones on the fibre channel switch or switches, if applicable.

Note: For information about zoning and enabling zones, see the documentation that is provided by the switch manufacturer.

4. Configure and verify the configuration.

FAST600:

The FAST600 Fibre Channel Storage Server only supports the use of 2 Gb fibre-channel switches. Set the port speed on the 2 Gb fibre-channel switches to 2 Gb.

Configuring JNI host bus adapter cards

The RDAC software includes the following shell scripts that you can use to configure JNI cards for use with connections through fibre channel switches:

- For FCI-1063 cards, `/etc/raid/bin/genfcaconf`
- For all `jnic.conf` cards, `/etc/raid/bin/genjnicconf`
- For FC64-1063 and all other cards, `/etc/raid/bin/genscsiconf`, which calls other scripts as appropriate

This section describes how to configure JNI HBA cards, including the following procedures:

- Installing the JNI adapter driver package
- Binding the HBAs to storage controllers (for SAN-attached configurations only)

Note: You must complete these procedures before you can install RDAC.

Steps for installing the JNI adapter driver package

This procedure refers to Table 18 on page 66.

Before you begin:

- Ensure that the JNI HBAs have been installed.
- Ensure that the JNI HBAs have been attached directly to the controllers or the JNI HBAs are attached to a SAN fabric switch.
- If attached to SAN fabric switch, the zones for the FAST Storage subsystem must be created and enabled. Consult the documentation provided by the switch manufacturer for information about how to create and enable zones.

Note: All fabric switches must be zoned in such a way that a single HBA can access only one controller per storage array.

Perform the following steps to install the JNI adapter driver package:

1. On Table 18 on page 66, identify the driver package for the installed HBAs.
2. Download the most current adapter driver package from the following Web site:

www.jni.com/Drivers

3. Run the following command to install the JNI adapter driver package:

```
#pkgadd -d adapter_driver_package
```

where *adapter_driver_package* is the name of the adapter driver package that you want to install, as in the following example:

```
#pkgadd -d JNIC146x.pkg
```

4. Run the following command to verify that the JNI adapter driver package is installed:

```
#pkginfo adapter_driver_package
```

where *adapter_driver_package* is the name of the adapter driver package that you installed.

Table 18 lists the configuration file for each JNI adapter driver package.

Table 18. JNI adapter configuration files

JNI HBA model numbers	JNI adapter driver package	Configuration file
FC64-1063-N	fcaw2.5.18.pkg	/kernel/drv/fcaw.conf
FCI-1063-N	fca-pci.2.5.18.pkg	/kernel/drv/fca-pci.conf
FCE2-1063, FCE2-1063, FCE-6410, FCE2-6412	JNIC.pkg, version 4.1.1.1	/kernel/drv/jnic.conf
FCE-1473, FCE-6412 (2 Gb), FCE-6460	JNIC146x.pkg, version 5.2.1	/kernel/drv/jnic146x.conf

Steps for modifying JNI HBA settings

The loop settings in the JNI configuration file are set by default to private loop for direct-attach configurations. For a SAN-attached configuration that uses a fabric switch, modify the loop settings to bind the HBAs to the FASTT storage controller before you install the RDAC driver on the host.

Attention: If you have a direct-attached configuration, skip this section and proceed to “Installing RDAC on Solaris hosts” on page 69. You do not need to modify JNI HBA settings if you have a direct-attached configuration.

This procedure refers to Appendix E, “JNI host bus adapter settings”, on page 147.

Before you begin: Ensure that the configuration file for your adapter meets the following criteria:

- The *FCLoopEnabled* and *FcFabricEnabled* variables must be set appropriately for your configuration; see the configuration file for details.
- If the *fca_nport* variable is in the file, set it to 1.
- If the *ip-disable* variable is in the file, set it to 1.
- In *fca*.conf* files, set *BusyRetryDelay* to 5000.
- In *fca*.conf* files, set *scsi_probe_delay* to 5000.
- In *fca*.conf* files, set *link_recovery_delay* to 1000.
- The lines that contain these variables must be uncommented and set to the specified values:
 - *failover* or *FailoverDelay*, set to 30
 - *JniCreationDelay*, set to 10
 - *def_wwnn_binding*, remove the dollar-sign character (\$) from the line
 - *def_wwpn_binding*
 - *def_hba_binding*

Notes:

1. *def_hba_binding* is for *fca.pci.conf* and *fcaw.conf*
2. When the RDAC software is installed, the *def_hba_binding* variable in the *fca.pci.conf* and *fcaw.conf* files is set to *nonjni*. This is the correct behavior.
3. *fcaw.conf* should be changed back to original value *-def_hba_binding="fcaw"*.

Perform the following procedures to modify JNI HBA settings:

1. Modify the loop attributes using the following steps:
 - a. Run the following commands:

```
#cd /kernel/drv  
#vi JNI_configuration_file
```

where *JNI_configuration_file* is the JNI configuration file whose HBA settings you want to modify, as in the following example:

```
#cd /kernel/drv  
#vi jnic146x.conf
```

- b. In the Vi editor, uncomment and modify the loop attributes using the information in Appendix E, “JNI host bus adapter settings”, on page 147.
- c. Run the following command to save changes made to the JNI configuration file:

```
#:wq
```

- d. Run the following command to restart the Solaris host:

```
#reboot -- -r
```

2. Gather binding information using the following steps:
 - a. After the host restarts, run the following commands to gather binding information in the messages log:

```
#cd /var/adm  
#more messages
```

The messages log appears.

- b. Search the messages log for the most recent HBA information. You can do this with a backwards search for the installed HBAs. For example, you can search for the JNI 2 Gb.HBA FCE-6064 by running the following command:

```
? FCE-6460
```

- c. After you find the most recent HBA information in the messages log, record the following information for future reference:
 - Target of the controller that is attached to the HBA
 - WWPN of the controller that is attached to the HBA
3. Bind the HBA to storage controllers using the following steps:

- a. Run the following commands to open the JNI configuration file:

```
#cd /kernel/drv
#vi JNI_configuration_file
```

For example, the following commands open the jnic146x.conf file:

```
#cd /kernel/drv
#vi jnic146x.conf
```

- b. Set and change the following variables to the specified values.

Note: The examples in this step assume you are binding two JNIC FCE-6460 HBAs to the FAStT controllers.

target0_hba 1. Set the target to the first record value from the /var/adm/messages file.

Example:

```
targetz_hba = "jnic146x0"
```

where z represents the target number for the first HBA from the /var/adm/messages file.

2. Set the target to the second record value from the /var/adm/messages file.

Example:

```
targety_hba = "jnic146x0"
```

where y represents the target number for the second HBA from the /var/adm/messages file

target_wwpn 1. Set the target to the wwpn of the first controller using the record value from the /var/adm/messages file.

Example:

```
targetz_wwpn="controller_1_wwpn"
```

where z represents the target number for the first HBA from the /var/adm/messages file.

2. Set the target to the wwpn of the second controller using the record value from the /var/adm/messages file.

Example:

```
targety_wwpn="controller_2_wwpn"
```

where y represents the target number for the second HBA from the /var/adm/messages file

- c. Run the following command to save the changes to the JNI configuration file:

```
#:wq
```

- d. Run the following command to restart the Solaris host:

```
#reboot -- -r
```

Installing RDAC on Solaris hosts

This section describes how to install RDAC.

Before you begin:

- You must install the JNI adapter driver package before you install RDAC. If you have a SAN-attached configuration, you must also set the bindings in the JNI configuration file before you install RDAC. If you fail to follow the procedures in this order, problems will occur.

For information about how to install the JNI adapter driver package, see “Steps for installing the JNI adapter driver package” on page 65.

For information about how to set the bindings, see “Steps for modifying JNI HBA settings” on page 66.

- Go to the following Web sites to ensure that you have the most recent RDAC:

```
ssddom02.storage.ibm.com/techup/webnav.nsf/support/fastt900
ssddom02.storage.ibm.com/techup/webnav.nsf/support/fastt600
ssddom02.storage.ibm.com/techup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techup/webnav.nsf/support/fastt500
ssddom02.storage.ibm.com/techup/webnav.nsf/support/fastt200
```

Perform the following steps to install RDAC:

1. Insert the Solaris installation CD in the CD-ROM drive.

Note: In this procedure, the installation CD is mounted at /cdrom/SM83. Adjust the procedure as required for your specific installation location.

2. Type the following command to start installing the RDAC package. :

```
# pkgadd -d /cdrom/SM83/SUN_Sparc/SM8rdac-Sparc-filename.pkg
```

The installation process begins.

Information about packages that can be installed in the specified directory is displayed on the command line, as in the following example:

```
The following packages are available:

1 RDAC                               Redundant Disk Array Controller
                                   (sparc) version number

Select package(s) you wish to process (or 'all' to process all
packages). (default:all) [?,??,q]:
```

3. Type the value of the package you are installing and press Enter.
The installation process begins.
4. The software automatically checks for package conflicts. If any conflicts are detected, a message is displayed indicating that some files are already installed and are in use by another package.

The following prompt is displayed:

```
Do you want to install these conflicting files [y, n, ?]
```

Type y and press Enter.

5. The following prompt is displayed:

```
This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <RDAC>

[y, n, ?]
```

Type y and press Enter.

The installation process continues.

6. When the RDAC package has been successfully installed, the following message is displayed:

```
Installation of <RDAC> was successful.
```

Ensure that the variables in the configuration files for your JNI adapter cards have been set to the correct values. For information about these configuration files and their variables, see “Steps for modifying JNI HBA settings” on page 66.

Attention: Modifying failover settings in the JNI configuration file after installing RDAC requires the removal of the RDAC from the host. In this case, perform the following procedure, in which you remove RDAC and restart the host.

Note: Modification or addition of targets or WWPNs for the controllers does not require the removal of RDAC or host restart. New HBAs can bind to controllers without the removal of RDAC or host restart. To modify or add targets or WWPNs for controllers, do not perform the following procedure; simply open and edit the JNI configuration file using the Vi editor.

If you need to modify the failover settings in the JNI configuration file, perform the following steps:

1. Uninstall the RDAC driver package using the following command:

```
#pkgrm RDAC_driver_pkg_name
```

2. Verify RDAC drive package removal using the following command:

```
#pkginfo RDAC_driver_pkg_name
```

3. Restart the Solaris host using the following command:

```
#reboot -- -r
```

4. Modify the JNI configuration file using the following command:

```
#cd /kernel/drv
#vi JNI_configuration_file
```

When you have finished making changes, run the following command to save changes:

```
#:wq
```

5. Install the RDAC driver package using the following command:

```
#pkgadd -d RDAC_driver_pkg_name
```

6. Verify package installation using the following command:

```
#pkginfo RDAC_driver_pkg_name
```

7. Restart the Solaris host using the following command:

```
#reboot -- -r
```

Verifying external storage

After you install RDAC, complete the following procedure to verify that the host can "see" your external storage.

Before you begin: SMutil must be installed on the host to enable the **SMdevices** and **hot_add** commands, which you use in the following procedure. For information about how to install SMutil, see "Installing SMutil" on page 60.

Note: Do not delete the access LUN or the access volume. RDAC uses the access LUN for communication between controllers and SMclient. Therefore, a storage subsystem with two controllers would require two access LUNs. The access LUN is typically assigned the LUN number 31, by default.

1. Run the following command to verify that your host can see your external FASTT storage:

```
#cd /etc/raid/bin/SMdevices
```

You must see two access LUNS, and all configured storage (logical drives). If any of these are missing, run the following command:

```
#cd /etc/raid/bin/hot_add
```

2. Run the following command to cross-verify between SMutils and what the host sees:

```
#format
```

You must see two access LUNS, and all configured storage (logical drives). If you still do not see all your storage check the following:

- Recheck your zones.
- Make sure the zones have been enabled.
- Recheck the targets and WWPNs
- Check to see the FASTT Controllers and expansion drawers are powered on.
- Check to see the switch(s) are powered on.
- Check fibre channel cables.

Redistributing volumes in case of failure

Auto Volume Transfer (AVT) is disabled, by default, on Solaris hosts. Therefore, if a failure occurs that initiates a controller failover, you must manually redistribute logical drives to their preferred paths.

Perform the following steps to redistribute logical drives to their preferred paths:

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage FAStT Installation and Support Guide* for the appropriate FAStT fibre-channel storage server.
2. Redistribute volumes to their preferred paths by clicking **Subsystem Management** —> **Storage Subsystem** —> **Redistribute Logical Drive**.

Chapter 5. Completing the software installation

This chapter describes how to complete the installation of Storage Manager 8.3 by performing these tasks:

- Failover protection
- Changing the host type
- Understanding the implications of using heterogeneous hosts in your storage subsystem
- Configuring the storage subsystem
- Configuring alert notifications

Failover protection

Failover protection is available by using multipath drivers. A multipath driver is an I/O path failover driver that is installed on host computers that access the storage subsystem. AVT/ADT is a built-in feature of the controller firmware that allows logical-drive-level failover rather than controller-level failover.

When AVT/ADT is enabled and used with a host multipath driver, it ensures that an I/O data path is available for the storage subsystem logical drives. The AVT/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 reestablishes ownership of the logical drive as soon as the multipath driver detects that the path is working again.

Changing the host type

The host type defines how the controllers in the storage subsystem work with the operating systems on the hosts that are connected to it.

Note: By default, AIX is the host type. If an incorrect host type is selected, existing volumes might be visible but redundancy through RDAC might fail.

When using the storage management software, you must set the correct host type by clicking **Storage Array** → **Change** → **Default Host Type** in the Subsystem Management window.

If partitioning is enabled, click **Mappings** → **Change** → **Host Type**. Select the desired host port in the left window and click **Configure** → **Topology** → **Change Host Type**.

For more information, see the topic on changing the host type of a storage subsystem in the Subsystem Management window online help.

If the Storage Partitioning feature is enabled, you must change the host type that is associated with each host port in the Mappings window. (For more information, see the topic on changing the host type of an individual host port in the Subsystem Management window online help.)

When you have selected the host type, continue with “Configuring storage subsystems” on page 74.

Understanding heterogeneous hosts

The heterogeneous hosts feature enables hosts that are running different operating systems to access a single storage subsystem. In previous releases of Storage Manager, only hosts running the same operating system could access a single storage subsystem. Storage Manager 8.3 supports up to 64 storage partitions, which enables a multiple host-type subsystem to share storage capacity, consolidate storage, and reduce storage management costs.

Host computers can be running completely different operating systems (for example, AIX and Solaris) or variants of the same operating system (for example, Solaris running in a cluster environment). When a host type is specified in the Define New Host Port window, the heterogeneous hosts feature enables the controllers in the storage subsystem to tailor their behavior (such as LUN reporting and error conditions) to the needs of the operating system or variant of the host that is sending the information.

In a heterogeneous environment, you must set each host type to the appropriate operating system during host-port definition so that the firmware on each controller can respond correctly for the operating system for that host. Before you begin setting up your heterogeneous host's configuration, see the *IBM FAStT Storage Manager Concepts Guide* available on your installation CD or at the following Web site:

www.ibm.com/pc/support/

Configuring storage subsystems

Use the following procedures to start the client software from either the storage management station or from a host that is acting as a storage management station (a host with the client software installed). Use the client software to configure each attached storage subsystem.

Attention: For cluster configurations, complete all applicable configuration procedures for each storage subsystem *before* installing the storage management software on a second host or cluster server.


Starting the Enterprise Management window

Use the following procedure to start the Enterprise Management window from the storage management station where you have installed the client software:

1. Start the client software:
 - If you are using a Windows workstation, click **Start** → **Programs** → **FAStT Storage Manager 8 Client**.
 - If you are using a UNIX[®]-based system, type `SMclient` at a shell prompt.

A splash screen is displayed while the client software starts. When the client software has been loaded, the Enterprise Management window and the Initial Automatic Discovery window opens (see Figure 11 on page 75).

Note: The Enterprise Management window can take several minutes to open. No wait cursor, such as an hourglass, is displayed.



SJ000704


Figure 11. Initial Automatic Discovery window

2. If you are running the client software from AIX, change to the /SMclient directory and then type SMclient.
3. If you are running the client software from Solaris 8, type SMclient.
4. Click **Yes** to begin an initial automatic discovery of all attached hosts and storage subsystems that are attached.

The software sends a broadcast message across the local subnetwork that is connected to the storage management station. It discovers host-agent-managed storage subsystems if the respective hosts respond to the broadcast. The software discovers directly managed storage subsystems if the controllers in the attached storage subsystems respond to the broadcast message.

It can take up to one minute for the Enterprise Management window to refresh after an initial automatic discovery. If you need to stop the automatic discovery operation for any reason, close the Enterprise Management window.

When the initial automatic discovery is finished, all attached hosts and attached storage subsystems are displayed in the Enterprise Management window (see Figure 12 on page 76).



SJ000705

Figure 12. Enterprise Management window

5. Ensure that all of the attached hosts and storage subsystems are displayed as expected. If not, do the following:
 - a. Check the hardware and connections for possible problems (see the hardware documentation for specific procedures).
 - b. See the Enterprise Management window help topic on discovering storage subsystems and take the appropriate action that is indicated.
 - c. Ensure that the device is on the local subnetwork. If it is not, click **Edit** → **Add Device** to add it. See the Enterprise Management window help topic on adding devices.

Note: Failure to add both controllers results in a partially-managed device.
 - d. If a storage subsystem is duplicated in the device tree after an automatic discovery, remove the duplicate storage subsystem icon from the device tree by clicking **Edit** → **Remove Device** in the Enterprise Management window.
6. Ensure that the status of each storage subsystem is **Optimal**:
 - If any device shows a status of **Unresponsive**, remove the device from the management domain, and then add it again. See the Enterprise Management window help topic on removing and adding devices.
 - If the device still shows an unresponsive status, contact your customer service representative.
7. Go to “Configuring alert notifications” on page 77.

Configuring alert notifications

After you have added devices to the management domain, you should set up alert notification options to report critical events on the storage subsystems. The following alert notification options are available:

- Alert notifications are sent to a designated network management station (NMS) using simple network management protocol (SNMP) traps.
- Alert notifications are sent to a designated e-mail address. See the Enterprise Management window help for specific procedures. To send e-mail to IBM, contact your customer service representative.
- Alert notifications are sent to a designated alphanumeric pager when third-party software is used to convert e-mail messages. See the Enterprise Management window help for specific procedures.

For more information about notification options, see *IBM FAStT Storage Manager Concepts Guide*, available on your installation CD or at the following Web site:

www.ibm.com/pc/support/

Use the following procedure to set up alert notifications using SNMP traps. You need to set up the designated management station only once.

1. Ensure that the installation CD is inserted in the CD-ROM drive on your designated NMS.
2. From the installation CD, copy the SM8.MIB file from the SM8mib directory to the NMS.
3. Follow the steps required by your NMS to compile the MIB. For details, contact your network administrator or see the documentation for the storage management product you are using.
4. Go to “Starting the Subsystem Management window”.

Starting the Subsystem Management window

In the Enterprise Management window, select a storage subsystem in one of these ways:

- Click **Tools** → **Manage Device**.
- Right-click the **storage subsystem** and click **Manage Device**.
- Double-click the **storage subsystem**.

The Subsystem Management window for the selected storage subsystem opens (see Figure 13 on page 78).

Note: You can manage one storage subsystem per Subsystem Management window. Open multiple windows to manage other storage subsystems.




Figure 13. Subsystem Management window

Renaming storage subsystems

When you start the storage management software for the first time, the storage subsystems are unnamed. Rename each storage subsystem from <unnamed> to a name that is appropriate for your network. In the Subsystem Management window, click **Storage Subsystem** → **Rename**.

Performing other storage subsystem management tasks

You can also perform other storage subsystem management tasks at this time. To perform the following tasks and other storage subsystem management tasks, see the appropriate topics in the Subsystem Management window help.

- Locate a storage subsystem
- View a storage subsystem profile
- Configure a storage subsystem password
- Create and manage arrays and array groups
- Use the performance monitor
- Create storage partitions (if applicable)

If you are installing the storage management software in a cluster environment, repeat all applicable client software installation procedures for a second storage management station and for all applicable host software on a second host. See Chapter 2, “Installing storage management station software on AIX systems”, on page 13, Chapter 3, “Installing storage management station software on HP-UX

systems”, on page 41, or Chapter 4, “Installing storage management station software on Solaris systems”, on page 55 for installation instructions.

Important

Create a storage subsystem profile and save it in a safe place whenever you modify the arrays and logical drives in your storage server. This profile contains detailed controller information including logical and physical disk configuration information. In the event of a catastrophic failure, it can be used to help recover the configuration.

To create a storage subsystem profile, select **View—> Storage Subsystem Profile** in the Storage Subsystem Management window and click the Save As button when the Storage Subsystem Profile window opens.

If you have a storage server with firmware 04.01.xx.xx or earlier, you must also perform controller state capture in addition to saving the storage subsystem profile. To use the Capture State Information feature, upgrade your controller to firmware 04.01.02.34 or later and install the IBM FASTT Storage Manager Field Tool, Version 5.21. This program is available at the following Web site:

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

Perform the following steps to create a storage subsystem profile using the Capture State Information feature:

1. Click **Tools—> Capture State Information**. A State Capture window opens.
2. Click the Browse button to specify the file name for the state capture information.
3. Click the Start button to start the capture process.
4. Click the Cancel button to close the State Capture window when the capture is complete.

Chapter 6. Storage Manager 8.3 with high-availability cluster services

The high-availability clustering services provided by Storage Manager 8.3 allow application services to continue when a hardware or software failure occurs. This system protects you from software failures as well as from the failure of a CPU, disk, or LAN component. If a component fails, its redundant partner component takes over cluster services and coordinates the transfer between components.

General information

Storage Manager 8.3 is certified for use with the following cluster services:

- MC/Service Guard versions A.11.09, A.11.13, and A11.14 on HP-UX systems
- Veritas Cluster Server 1.3, 2.0, and 3.5 on Solaris systems
- Veritas Volume Manager 3.2, and 3.5 on Solaris systems.
- Veritas File Systems 3.2 and 3.5 on Solaris systems
- HACMP ES and ESCRM 4.4.1 on AIX systems
- HACMP ES and ESCRM 4.5 on AIX systems
- PSSP 3.4 with GPFS 1.5 on AIX systems
- PSSP 3.5 with GPFS 2.1 on AIX systems

This document does not describe how to install or configure cluster services. Refer to documentation provided with your cluster service products for this information.

Note: When using storage partitioning in conjunction with a cluster solution, all partitions in the clusters must "see" the same storage. Therefore, logical drive mapping to cluster partitions must be done on the host group level.

FAST600:

HACMP, PSSP and GPFS are not supported on the FAST600 Fibre Channel Storage Server.

Prerequisites for HP-UX

You can choose among many configurations when you set up clustering on an HP-UX system. A minimum configuration consists of two servers that are configured with both a primary and two standby LANs to establish a heartbeat LAN.

Provide fibre connections to the storage subsystem through two switches that provide the necessary redundant data path for the hosts. Ensure that each server has two HP Tachyon host bus adapters.

Prerequisites for Solaris and Veritas Cluster Server

The following sections contain general hardware requirements and additional information about the cluster services.

General hardware requirements

Each Solaris system in the cluster requires the following hardware:

- At least three Ethernet ports:
 - Two for the private network connections
 - At least one for the public network connection
- Two fibre host bus adapters for connection to the storage subsystem
- A SCSI connection for operating system disks
- Each Veritas Cluster Server system requires at least 128 MB of RAM and 35 MB of free disk space

System dependencies

This section provides information about Veritas Volume Manager, RDAC IDs, and single points of failure.

Veritas Volume Manager and Data Multipathing

If you are using the Veritas Volume Manager, you must disable Data Multipathing (DMP) either for all devices, or for just FASTT storage servers, depending on your version of the Veritas Volume Manager.

For Veritas Volume Manager 3.1 or earlier, follow these steps to disable DMP for all devices:

1. Open the `/etc/system` file in the vi text editor by typing the following command:

```
# vi /etc/system
```

2. Comment out the `forceload: drv/vxdmp` line.
3. Save and close the `/etc/system` file.

For Veritas Volume Manager 3.2 and later, you must disable DMP for FASTT storage servers only. Use option 17 and then option 5 of the submenu of the `vxdiskadm` utility to do so. Other devices can use DMP.

Note: If you are using the Veritas Cluster Server, you must change the stack size parameters. Open the `/etc/system` file in the vi test editor and make the following updates:

- Change the default value for `set lwp_default_stksize` to `0x8000`
- Change the default value for `set rcpcmod:svc_default_stksize` to `0x8000`

RDAC IDs

Add up to eight additional IDs to the `/etc/symsm/rmparams` file. Complete the following steps to add them:

1. Open the `/etc/symsm/rmparams` file in the vi text editor by typing the following command:

```
# vi /etc/symsm/rmparams
```

2. Modify the `Rdac_HotAddIDs` line as follows:

```
Rdac_HotAddIDs:0:1:2:3:4:5:6:7:8
```

3. Save and close the `/etc/symsm/rmparams` file.

Single points of failure

When setting up cluster services, it is important to eliminate single points of failure because a single point of failure makes a cluster only as strong as its weakest component. Set up the storage subsystem for shared storage; for example, all the nodes in the cluster must recognize the same storage and the host types must be set correctly.

Storage Manager 8.3 with High Availability Cluster Management Program (HACMP)

The following section contains general requirements and usage notes for HACMP support with Storage Manager 8.3.

FAStT600: HACMP is not supported on the FAStT600 Fibre Channel Storage Server.

FAStT software and firmware requirements

For hardware, software, firmware and RDAC requirements, see Appendix A, “AIX system requirements”, on page 139.

For the most current information on AIX and Device Driver versions for FAStT, see the readme.txt file posted at the following Web site:

www.ibm.com/pc/qtechinfo/MIGR-50177.html

After accepting the license agreement, proceed to the v8.3 link for AIX.

HACMP software requirements

The following HACMP software requirements are in addition to requirements specified in Appendix A, “AIX system requirements”, on page 139.

Table 19. HACMP software requirements

HACMP and AIX versions	Required HACMP APARs	Required AIX APARs	
HACMP ES and ESCRM 4.4.1 with AIX 4.3.3 and AIX 5.1.0	IY31483 IY31489 IY38234	None	
HACMP ES and ESCRM 4.5 with AIX 5.1.0 and AIX 5.2.0	IY32874 IY37247	AIX 5.1: IY37897	AIX 5.2: IY37746 IY36626

Configuration limitations

The following limitations apply to HACMP configurations:

- Switched fabric connection only; no direct connection is allowed between the host node and FAStT.
- HACMP C-SPOC cannot be used to add a FAStT disk to AIX using the *Add a Disk to the Cluster* facility.
- HACMP C-SPOC does not support enhanced concurrent mode volume groups.
- Single node quorum is not supported in a two-node GPFS cluster with FAStT disks in the configuration.
- HACMP is not supported in a heterogeneous environment.

- HACMP is not supported on FAStT600.

Other HACMP usage notes

The following notations are specific to HACMP environments:

- Concurrent and non-concurrent modes are supported with HACMP 4.4.1 and 4.5 and FAStT running Storage Manager 8.3, including hot standby and mutual takeover.
- HACMP 4.4.1 and 4.5 are supported on Regatta LPAR clustered configurations.
- HACMP clusters can support from two to 32 servers on each FAStT partition. If you run this kind of environment, be sure to read and understand the AIX device drivers queue depth settings that are described in Chapter 2, “Installing storage management station software on AIX systems”, on page 13.
- Non-clustered AIX hosts can be connected to the same FAStT that is attached to an HACMP cluster with HACMP version 4.4.1 and version 4.5, and FAStT running Storage Manager 8.3. However, the non-clustered AIX hosts must be configured on separate FAStT host partitions.

Storage Manager 8.3 with Parallel System Support Programs (PSSP) and General Parallel File System (GPFS)

The following section contains general requirements and usage notes for Parallel System Support Programs (PSSP) and General Parallel File System (GPFS) support with Storage Manager 8.3.

FAStT600:

PSSP and GPFS are not supported on the FAStT600 Fibre Channel Storage Server.

PSSP, and GPFS software requirements

The following combinations of PSSP and GPFS are supported with FAStT500, FAStT700 and FAStT900:

- PSSP 3.4 with GPFS 1.5
- PSSP 3.5 with GPFS 2.1

Configuration limitations

The following limitations apply to PSSP and GPFS configurations:

- The FAStT200 and FAStT600 are not supported.
- Direct connection is not allowed between the host node and a FAStT storage server. Only switched fabric connection is allowed.
- Each server must have two or four fibre-channel connections to each attached FAStT storage server; one connection for each FAStT controller host-side minihub; and two connections for each FAStT partition with each connection on different controllers. These connections must be made using two or four Fibre Channel host adapters. Multiple FAStT storage servers can be attached using these same host adapters (fan out) or by adding additional HBA pairs for each partition.
- A maximum of four partitions are allowed for each FAStT storage server.
- RVSD clusters can support up to two IBM Virtual Shared Disk and RVSD servers for each FAStT partition.

- HACMP/GPFS clusters can support from two to 32 servers on each FAStT partition. If you run this kind of environment, be sure to read and understand the AIX device drivers queue depth settings that are described in Chapter 2, “Installing storage management station software on AIX systems”, on page 13.
- For maximum availability, distribute the HBA and FAStT connections across separate fibre-channel switches to minimize the effects of a SAN fabric failure.
- Single node quorum is not supported in a dual-node GPFS cluster with FAStT disks in the configuration.

Other usage notes specific to PSSP/GPFS environments

In GPFS file systems, the following FAStT cache settings are supported:

- Read cache enabled or disabled
- Write cache enabled or disabled
- Cache mirroring disabled

The performance benefits of read or write caching depends upon the application.

GPFS/PSSP logical configuration diagrams

The diagrams in this section show both the preferred and failover paths from an HBA pair to a given volume or set of volumes.

A preferred path to a volume is determined when the volume is created and distributed across a FAStT controller. The controller to which it is assigned determines which path is preferred or active for I/O transfer. Volumes can, and in most cases should, be assigned to both controllers, balancing the I/O load across HBAs and FAStT controllers.




Figure 14. RVSD cluster configuration with single FASt500/700/900 - Storage Manager 8.3, shown with two partitions (four maximum) per FASt




Figure 15. RVSD cluster configuration with multiple FAST500/700/900 - Storage Manager 8.3, shown with one partition per FAST




Figure 16. Typical RVSD cluster configuration with multiple FAS T500/700/900 - Storage Manager 8.3, shown with one partition per FAS T




Figure 17. RVSD cluster configuration with multiple FAST500/700/900 - Storage Manager 8.3, shown with two partitions per FAST




Figure 18. HACMP/GPFS cluster configuration with single FAST500/700/900, shown with one partition per FAST




Figure 19. HACMP/GPFS cluster configuration with multiple FAST500/700/900, shown with two partitions per FAST

Chapter 7. Copy Services and the FAS*t*T Storage Server

FAS*t*T Storage Manager 8.3 supports the following premium features, which are available for purchase separately from IBM or an IBM Business Partner:

FlashCopy

The FlashCopy premium feature supports the creation and management 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.

Remote Mirror Option

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

This document describes the FlashCopy premium feature.

For Remote Mirror Option concepts, configuration procedures, and operation and administration tasks, see *IBM FAS*t*T Remote Mirror Option Installation and User's Guide*.

FAS*t*T600:

FlashCopy and Remote Mirror Option are not supported on the FAS*t*T600 Fibre Channel Storage Server.

Overview of FlashCopy

FlashCopy supports the creation and management of *FlashCopy logical drives*. A FlashCopy logical drive is a point-in-time (PIT) 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.

You can create more than one FlashCopy 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.

This chapter describes the following procedures, which you perform using the SMclient:

- How to enable FlashCopy
- How to plan a FlashCopy logical drive
- How to create a FlashCopy logical drive

- How to perform FlashCopy tasks

You can also perform these procedures using the command-line interface (CLI) and the script editor. For more information, click on the **Help** tab on the Enterprise management window, and then search for *SMcli* or *script editor*.

Enabling FlashCopy

To enable FlashCopy, perform the procedures that are described in this section in the following order:

1. Obtaining the feature enable identifier
2. Generating the feature key file and enabling FlashCopy

Before you begin: Gather the following information:

- Your Feature Activation Code, which accompanies the FlashCopy Premium Feature
- Your controller unit IBM serial number, which is printed on the label on the back of your FAStT controller unit

Ensure that your controller unit and expansion units are connected, powered on and configured.

Obtaining the feature enable identifier

Each storage subsystem has its own unique feature enable identifier. This identifier ensures that a particular feature key file is applicable only to that storage subsystem. Complete the following steps to obtain the feature enable identifier:

1. Click **Start** → **Programs** → **FAStT Storage Manager 8.3 Client**. The Enterprise Management window opens.
2. In the left-hand pane, double-click the storage subsystem for which you want to enable FlashCopy. The Subsystem Management window opens for the selected storage subsystem.
3. In the Subsystem Management window, click **Storage Subsystem** → **Premium Features** → **List**. The Premium Feature List window opens and displays the feature enable identifier.
4. Record the feature enable identifier.
5. Close the Premium Feature List window.

Generating the feature key file and enabling FlashCopy

You enable FlashCopy feature using a tool that is located at the following Web site:

www.ibm.com/storage/fasttkeys

1. Generate and store the key activation file by completing Step 1 on the Web site.
2. Download the key activation file to your controller unit by completing Step 2 on the Web site.
3. Enable FlashCopy by completing Step 3 on the Web site.
4. Verify that FlashCopy was successfully enabled by completing Step 4 on the Web site.

Your system is now ready to use FlashCopy.

If FlashCopy becomes disabled, you can access the Web site and repeat this process. For any additional assistance, please contact your local IBM service provider.

Planning a FlashCopy logical drive

When you create a FlashCopy logical drive, you specify where to create the FlashCopy repository logical drive. In addition, you specify its capacity, threshold level warning, and other parameters. Before you create a FlashCopy logical drive, be familiar with the following concepts and tasks in this section:

- Setting FlashCopy repository logical drive capacity options
- Viewing current FlashCopy repository logical drive capacity settings
- Estimating FlashCopy repository logical drive capacity
- Estimating FlashCopy repository life

Setting FlashCopy repository logical drive failure options

The FlashCopy repository logical drive capacity is created as a percentage of the base logical drive. It contains the copy-on-write data.

The default FlashCopy repository logical drive capacity is initially set to 20% of the base logical drive, if enough free capacity exists to create a FlashCopy repository logical drive of this size. The default threshold level for the FlashCopy repository logical drive is initially set to 50%.

If you are not sure how large to make the FlashCopy repository logical drive or how high to set the *FlashCopy repository logical drive full* warning, accept the default settings. You can estimate later how quickly the repository capacity is being used. For more information, see “Estimating FlashCopy repository life” on page 99.

The storage management software provides a warning message when your FlashCopy repository logical drive exceeds the threshold level. Do not ignore this notification. It is the only warning that you receive before the FlashCopy repository logical drive becomes full. When you receive the warning, you can perform one of the following two actions:


- Increase the capacity of the FlashCopy repository logical drive
- Increase the FlashCopy repository logical drive threshold capacity warning level

If you choose to increase the threshold capacity warning level, you reduce the amount of time that you have to respond the next time the threshold is reached and you receive a *threshold exceeded* notification.

Viewing current FlashCopy repository logical drive failure settings

Complete the following steps to view the current FlashCopy logical drive failure settings.


1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** —> **Properties**, or right-click and select **Properties**. The FlashCopy Repository Logical Drive - Properties window opens. See Figure 20 on page 96..



SJ000825

Figure 20. FlashCopy repository logical drive properties

3. Click the **Capacity** tab to view the currently defined settings, as shown in Figure 21 on page 97.



SJ000826

Figure 21. Repository capacity settings

If the repository full policy is set to *Fail FlashCopy logical drive* (default setting), the data is not recoverable when the drive becomes full and FlashCopy cannot be accessed. If this situation occurs, you can perform one of the following actions:

- Delete the FlashCopy logical drive
- re-create the FlashCopy logical drive to create a new point-in-time image

If the repository full policy is set to *Fail writes to base logical drive*, the data is recoverable. But, the FlashCopy repository logical drive capacity must be increased before writes to the base logical drive are not rejected.

Attention: The *Fail writes to base logical drive* option of the repository full policy is not supported on AIX. Selecting this option might cause data loss on the base logical drive. If you are using AIX, ensure that default option, *Fail FlashCopy logical drive* is selected. Be sure to monitor the capacity of the FlashCopy repository logical drive, because the you cannot access the FlashCopy if the repository logical drive becomes full.

For more information, see “Resizing a FlashCopy repository logical drive” on page 134.

Note the following points:

- Deleting a FlashCopy logical drive automatically deletes the associated FlashCopy repository logical drive.
- Deleting a FlashCopy repository logical drive automatically deletes the associated FlashCopy logical drive.

- Deleting a FlashCopy logical drive and then creating it again forces you to stop the host application and unmount the base logical drive while the FlashCopy is being created again.
- Recreating a FlashCopy logical drive alleviates the need to create a FlashCopy repository logical drive, as well as remap the assigned logical drive-to-LUN mappings between the FlashCopy logical drive and the host.
- After you re-create the FlashCopy logical drive, you can change parameters on the FlashCopy repository logical drive through the appropriate menu options.
- To avoid another *FlashCopy repository logical drive full* failure, increase the capacity of the FlashCopy repository logical drive. For more information, see “Resizing a FlashCopy repository logical drive” on page 134.

Estimating FlashCopy repository logical drive capacity

When a FlashCopy logical drive is created, a physical logical drive called the *FlashCopy repository logical drive* is also created to store FlashCopy data and copy-on-write data. The default setting for the FlashCopy repository logical drive capacity is 20% of the base logical drive’s capacity. In general, this capacity should be sufficient. Use the following information to help determine the appropriate capacity of the FlashCopy repository logical drive:

- A FlashCopy repository logical drive can be no smaller than 8 MB.
- The amount of write activity to the base logical drive after the FlashCopy logical drive has been created dictates the size of a FlashCopy repository logical drive. As the amount of write activity to the base logical drive increases, the number of original data blocks that need to be copied from the base logical drive to the FlashCopy repository logical drive also increases.
- The estimated life expectancy of the FlashCopy logical drive contributes to determining the capacity of the FlashCopy repository logical drive. The longer a FlashCopy logical drive remains enabled, the FlashCopy repository logical drive runs the increased risk of reaching its maximum capacity. For more information, see “Estimating FlashCopy repository life” on page 99.
- The amount of management overhead required on the FlashCopy repository logical drive to store FlashCopy logical drive data contributes to determining the capacity of the FlashCopy repository logical drive. The amount of management overhead actually required is fairly small and can be calculated using the simple formulas detailed in the following section.
- There is not necessarily a one-to-one correlation between the number of data blocks that change on the base logical drive and the amount of copy-on-write data that is stored on the FlashCopy repository logical drive. Depending on the location of data blocks that need to be copied, for performance reasons, the controller might copy over a full set of 32 blocks, even if only one set of blocks has changed. Keep this in mind when determining the percentage of the capacity of the base logical drive that can be copied to the FlashCopy repository logical drive.

Calculating expected overhead

Use the following formula to calculate the amount of management overhead that is required to store FlashCopy data on the FlashCopy repository logical drive. This formula should be used merely as a guide, and FlashCopy repository logical drive capacity should be re-estimated periodically.

Note: For this formula, you must convert the bytes to kilobytes, and then convert the kilobytes to megabytes.

The formula to calculate the amount of management overhead required is:

$$192 \text{ KB} + (X / 2000)$$

Where X is the capacity of the base logical drive in bytes.

Example

For a 5 GB base logical drive, where 30% of the data blocks are expected to change on the base logical drive, the estimated FlashCopy repository logical drive capacity can be calculated as follows:

1. Convert the capacity of the base logical drive to bytes.
When converted, 5 GB equals 5 368 709 120 bytes.
2. Divide the capacity of the base logical drive (in bytes) by 2000.
When divided, the result is 2 684 354.56 bytes.
3. Convert the result from step 2 (in bytes) to kilobytes (KB).
When converted, the result is 2621.44 KB.
4. Add 192 KB to the results from step 3.
 $192 \text{ KB} + 2621.44 \text{ KB} = 2813.44 \text{ KB}$
5. Convert the result from step 4 to megabytes (MB).
When converted, the amount of management overhead required is 2.75 MB (or 0.002686 GB).
6. In this example, 30% of the data blocks on the base logical drive are expected to change. To accurately calculate the FlashCopy repository logical drive capacity, sufficient space needs to be allocated for the copy-on-write data as well as the management overhead (calculated in step 5).
To calculate the copy-on-write space required, calculate the percentage of the base logical drive expected change:
 $30\% \text{ of } 5 \text{ GB} = 1.5 \text{ GB}$
The final estimated FlashCopy repository logical drive capacity for this example is:
 $1.5 \text{ GB} + 0.002686 \text{ GB} = 1.502686 \text{ GB}$
7. In the Create FlashCopy Logical Drive wizard: Specify Repository Capacity window, use the **percent (%) full** box field to set the estimated FlashCopy repository logical drive capacity. See Figure 21 on page 97.

Note: The **percent (%) full** field sets the FlashCopy repository logical drive capacity as a percentage of the base logical drive. Using the **percent (%) full** field, increase or decrease the percentage until the FlashCopy Repository Logical Drive Capacity value matches the estimated capacity calculated in step 6. (Some rounding up might be required.)

Estimating FlashCopy repository life

When you create a FlashCopy logical drive, you are asked to define various properties for the FlashCopy repository logical drive, including the name of the FlashCopy repository logical drive, its capacity, a logical drive-to-LUN mapping, and the repository full condition. When defining the properties of the FlashCopy repository logical drive, keep in mind the kind of usage you have planned for the FlashCopy logical drive. Understanding how the FlashCopy logical drive will be used can help you to estimate the life expectancy of the FlashCopy repository logical drive.

If numerous I/O requests are written to the base logical drive, the FlashCopy repository logical drive (containing the FlashCopy data and copy-on-write data) could eventually exceed the base logical drive capacity if all the original data blocks

are changed. The default setting suggests 20% of the base logical drive capacity, but this is a setting that can be adjusted after some usage data becomes available.

The following procedure describes how to estimate the life expectancy of a FlashCopy repository logical drive. Use it merely as a guide. Note that the life expectancy of a FlashCopy repository logical drive should be carefully re-estimated periodically.

1. Select the FlashCopy repository logical drive in the Logical view. Click **Logical Drive** → **Properties** and click the **Capacity** tab. Alternately, you can right-click the FlashCopy repository logical drive, select **Properties**, and click the **Capacity** tab.
2. Select the FlashCopy repository logical drive in the Logical view. Click **View** → **GoTo** → **FlashCopy Logical Drive**. Alternately, right-click the FlashCopy repository logical drive and select **Go To FlashCopy Logical Drive**.
3. Record the creation timestamp day and time.
4. Record the capacity used (GB) and the available capacity (GB).
5. Determine the elapsed time (t) by subtracting the creation time from the current time, expressing the elapsed time in either minutes, hours, or days.
6. The total time (Tr) that the repository is available for copy-on-write data can now be estimated (based on the current usage) by multiplying the elapsed time (t) by the available capacity (Ct), and then dividing the resultant number by the capacity used (Cu).

The formula you need to use is:

$$Tr = (t * Ct) / Cu$$

Where:

Tr	Total time available
t	Elapsed time
Ct	Available capacity
Cu	Capacity used

Note: The total time available (Tr) indicates the total usage time for the FlashCopy repository logical drive.

Using the result of these calculations, you should now be able to make an informed decision about whether you should increase the capacity of the FlashCopy repository logical drive. If the repository capacity becomes 100% full during the expected lifetime of the FlashCopy, then you should increase the capacity of the FlashCopy repository logical drive. For more information, see “Resizing a FlashCopy repository logical drive” on page 134.

Creating a FlashCopy logical drive

This section contains general information and procedures for how to create a FlashCopy logical drive from a base volume.

See “Creating FlashCopy logical drives on AIX, HP-UX, and Solaris” on page 110 for the following procedures that are specific to your operating system:

- How to create a FlashCopy logical drive on AIX
- How to create a FlashCopy logical drive on HP-UX
- How to create a FlashCopy logical drive on Solaris

Note:


- See the Create a FlashCopy Logical Drive wizard online help for additional instructions for the operating-system-specific procedures. Failure to complete the additional steps that are required for your host operating system can result in an inaccurate point-in-time image of the base logical drive.
- You cannot create a FlashCopy logical drive of a base logical drive that is the secondary logical drive in a Remote Volume Mirror.

Before you begin: Read the information in “Using the script editor and command-line interface” on page 121.

Perform the following steps to create a FlashCopy logical drive:

1. Place the application into backup mode:
 - a. Stop the host application that is accessing the base logical drive, sync file system. Unmount the base logical drive if possible. This enables a valid consistent copy to be taken.

It might not be convenient to stop database applications. However, in this case, it is required to place the application into a backup mode or place it in an acquiesced state while you create the FlashCopy.
 - b. It is important to back up application recovery files, such as role back and redo logs, because these files might be located in different physical disk storage or on different logical drives.
2. Launch the Create a FlashCopy Logical Drive wizard:
 - a. Select a base logical drive from the Logical view.
 - b. Click **Logical Drive** → **FlashCopy** → **Create**. Alternately, right-click and select **Create FlashCopy Logical Drive**. See Figure 22.



SJ000832

Figure 22. Create FlashCopy menu options

The Create a FlashCopy Logical Drive wizard opens as shown in Figure 23.





Figure 23. Create a FlashCopy Logical Drive Wizard startup

Note: If the FlashCopy logical drive is to be based on the root disk of the host operating system, the final point-in-time image might not be completely consistent with the base logical drive.

3. Create the FlashCopy logical drive:
 - a. Review the information in the initial window as shown in Figure 23. Click **OK** to proceed to the wizard introduction window.
Follow the instructions on each wizard panel, and click the **Next** button when you are ready to continue to the next panel.

Note: Each wizard panel has context-sensitive help. Click the **Help** button on a particular panel to receive help for that panel.


- b. The Introduction window (Figure 24 on page 103) defines what a FlashCopy logical drive is and the physical components that are associated with a FlashCopy logical drive. It enables you to select either the Simple or Advanced path through the Create FlashCopy Logical Drive wizard:
 - **Simple Path:** Proceeds to the **Specify Name** panel, which provides a preview of the FlashCopy and repository default names. You can also change the defaults on this panel.
 - **Advanced Path:** Proceeds to the **Allocate Capacity** panel. There you select the Free Capacity or Unconfigured Capacity node on which to place the FlashCopy repository logical drive.
If no free capacity exists or the available free capacity is unusable, a warning message displays.



SJ000834

Figure 24. Wizard Introduction window


- c. If you want to use the Advanced path, select **Advanced** and click **Next**. The Allocate Capacity window (Figure 25 on page 104) opens. From this window you can choose free capacity or unconfigured capacity. A description of each option follows.



SJ000835

Figure 25. Allocate Capacity window

- **Free Capacity:**
If you select **Free Capacity** and click **Next**, the Specify Logical Drive Parameters window (Figure 26 on page 105) opens.



SJ000836

Figure 26. Specify Logical Drive Parameters

Define the *FlashCopy logical drive-to-LUN mapping* parameter and the *FlashCopy repository logical drive full policy*.


Attention: The *Fail writes to base logical drive* option of the *FlashCopy repository logical drive full policy* is not supported on AIX. Selecting this option might cause data loss on the base logical drive. If you are using AIX, select the *Fail FlashCopy logical drive* option, which is the default.

Click **Next**.

The Preview window opens and displays the FlashCopy and repository parameters. You can click **Back** to return to the previous windows and edit the parameters or click **Finish** and continue.

- **Unconfigured Capacity:**

In the Allocate Capacity window (Figure 25 on page 104), if you select **Unconfigured Capacity** and click **Next**, the Specify Array Parameters window (Figure 27 on page 106) opens. Here you can specify a new array on which the repository resides. You can also specify the RAID level of the array that meets the FlashCopy repository logical drive data storage and protection requirements.



SJ000838

Figure 27. Specify Array Parameters window

After you select where you are placing the volume, same array, and free capacity on another array, or you create a new array, the steps for defining the FlashCopy logical drive parameters are the same as for the Free Capacity option.

On the Specify Logical Drive Parameters window (Figure 26 on page 105), define the FlashCopy logical drive-to-LUN mapping parameter, the threshold percentage full parameter, and the FlashCopy repository logical drive full conditions.

The Preview window provides a preview of the FlashCopy and repository parameters. You can click **Back** and return to the previous windows to edit the repository parameters, or click **Finish** and continue.


- d. On the Specify Names window (Figure 28 on page 107), define the FlashCopy logical drive name and the name of its associated FlashCopy repository logical drive.

The default naming convention for the first FlashCopy uses the base volume name and adds a suffix of "-1" for the logical drive and "-R1" for the repository drive. The second FlashCopy increments this number to 2, and so on, up to the four volumes.

For example, if you are creating the first FlashCopy logical drive for a base logical drive called *DataVol*, then the default FlashCopy logical drive name is *DataVol-1*, and the associated FlashCopy repository logical drive default name is *DataVol-R1*. The default name of the next FlashCopy logical drive you create based on *DataVol* is *DataVol-2*, with the corresponding FlashCopy repository logical drive named *DataVol-R2*, by default.

Change the default names if required.

Click **Next** to continue.




SJ000839

Figure 28. Specify Names window

Tips about names:

- Regardless of whether you use the software-supplied sequence number that (by default) populates the **FlashCopy logical drive name** or **FlashCopy repository logical drive name** field, the next default name for a FlashCopy or FlashCopy repository logical drive still receives a sequence number determined by the software. For example, you might name the first FlashCopy of base logical drive DataVol *DataVolMay28*, without using the software-supplied sequence number of 1. Regardless, the software assigns a default name for the next FlashCopy as *DataVol-2*.
- The next available sequence number is based on the number of existing FlashCopies of a base logical drive. If you delete a FlashCopy logical drive, its sequence number becomes available again.
- You must choose a unique name for the FlashCopy and FlashCopy repository logical drives. Otherwise, the system displays an error message.
- Names have a 30-character limit. When you reach this limit in either the **FlashCopy logical drive name** or **FlashCopy repository logical drive name** field, it no longer accepts input. If the base logical drive name contains 30 characters, the default names for the FlashCopy and its associated FlashCopy repository logical drive use the base logical drive name truncated just enough to add the sequence string. For example, for “Host Software Engineering Group GR-1”, the default FlashCopy name would be “Host Software Engineering GR-1”. The default repository name would be “Host Software Engineering G-R1.”

- e. On the Specify FlashCopy Repository Logical Drive's Capacity window (Figure 29), set the repository drive capacity as a percentage of the capacity of the base logical drive.



SJ000840

Figure 29. Specify Repository Drive Capacity


Use the **percent (%) of base logical drive** field to set the desired capacity. Click **Next** to continue.

The capacity requirement varies, depending on the frequency and size of I/O writes to the base logical drive and the length of time that you need to retain the FlashCopy logical drive. In general, you should choose a larger capacity for the repository if you intend to keep the FlashCopy logical drive for a long period of time. You should choose a larger capacity if a large percentage of data blocks will change on the base logical drive during the life of the FlashCopy logical drive due to a high level of I/O activity. Use historical performance monitor data or other operating system utilities to help you determine typical I/O activity levels on the base logical drive.

Note: In most situations, the 20% default value should be sufficient capacity for your FlashCopy repository logical drive. For information about determining the optimal size, see “Estimating FlashCopy repository logical drive capacity” on page 98, and “Estimating FlashCopy repository life” on page 99.

- f. The Create FlashCopy Logical Drive wizard - Preview window opens. It displays components that are associated with the FlashCopy. Review the information and click **Next**.


- g. The Create FlashCopy Logical Drive wizard - Completed window (Figure 30) opens. It displays the associated logical drives and components that make up the FlashCopy relationship. Click **Finish** to continue.



SJ000841

Figure 30. Completed window

- h. The FlashCopy drives are now displayed in the device management GUI, as shown in Figure 31 on page 110. See "Viewing the FlashCopy icon" on page 123 for information about interpreting the display icons.



SJ000842

Figure 31. FlashCopy volumes

- Restart the host application. After you create one or more FlashCopy logical drives, mount the base logical drive and restart the host application using that base logical drive.

Creating FlashCopy logical drives on AIX, HP-UX, and Solaris

This section explains the specific procedures for creating FlashCopy logical drives on UNIX-based hosts. It outlines the key steps that are required to create and reuse a FlashCopy logical drive on the following operating systems:

- AIX
- Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8)
- HP-UX 11.0 (or higher)

Instructions for AIX: Logical Volume Manager

Use the following procedure to create FlashCopy logical drives on a host running AIX, using Logical Volume Manager (LVM).

You can create FlashCopy logical drives for repeated reuse (frequent or nightly backups) or for one-time use (speculative change or upgrade testing).

Failure to complete the steps listed in the following procedure can result in an inaccurate point-in-time image of the base logical drive.

Creating a FlashCopy logical drive on AIX:

Before you begin:

- Read the information in “Using the script editor and command-line interface” on page 121.
- Stop all I/O activity to the base logical drive, or suspend data transfer. This ensures that an accurate point-in-time image of the base logical drive is captured.

Restrictions:

- FlashCopy logical drives can be created only for AIX arrays. If an array has more than one logical drive, you must create a FlashCopy logical drive for each logical drive in the array.
- AIX does not support the *Fail writes to base logical drive* option of the repository full policy. Selecting this option might cause data loss on the base logical drive. You must ensure that the repository full policy is set to the default option, *Fail FlashCopy logical drive*.

For information about how to set the repository full policy, see “FlashCopy repository properties” on page 127

Perform the following steps to create a FlashCopy logical drive:

1. Locate the array on which the FlashCopy logical drive is to be based, and unmount each of its associated filesystems using the following command:

```
umount mount-point
```

where *mount-point* is the name of the filesystem that is being unmounted.

2. Start the Storage Management software. The Enterprise Management window opens.
3. Launch a Subsystem Management window, using one of the following methods:
 - Select the storage subsystem in the Device Tree view or from the Device Table. Then, click the **Manage Device** toolbar button, or click **Tools** → **Manage Device**.
 - Right-click the storage subsystem in the Device Tree view or Device Table and select **Manage Device**.
 - Double-click a storage subsystem node in the Device Table.
 - Select the storage subsystem in either the Device Tree view or from the Device Table and press Enter.

The Subsystem Management window opens in a separate window.

4. Within the Logical view of the Subsystem Management window, select a standard logical drive, and create a FlashCopy logical drive using the Create a FlashCopy Logical Drive wizard. This is accessed from the Subsystem Management window.

Important: If an AIX array has more than one logical drive, you must create a FlashCopy logical drive for each logical drive in the array.
5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive. Logical drive-to-LUN mappings can be defined using the Storage Partitioning wizard to quickly define a single storage partition. It guides you through the major steps to specify which host will access a logical drive, and the associated LUNs.
6. Log in to the host as root.
7. Ensure that the host operating system recognizes the FlashCopy logical drives. At the host prompt, type the following command and press Enter.

```
cfgmgr
```

Several minutes might pass while the operating system accesses the drives. When the operation completes, a window opens with the following message:

```
Device nodes have been updated
```

The new logical drives are now available from the operating system.

8. At the host prompt, type the following command and press Enter.

```
lspv
```

A list of the physical drives recognized by the operating system is displayed.

9. Look for the operating system device name of your FlashCopy logical drive in the list. The list shows a physical logical drive ID (PVID) for the FlashCopy logical drive. This is the same as the PVID for the associated base logical drive, because the FlashCopy logical drive contains the same array data structures as the base logical drive.
10. Clear the PVID for the FlashCopy logical drives. At the host prompt, type the following command and press Enter.

```
chdev -l os_device_name -a pv=clear
```

where *os_device_name* is the operating system device name of the FlashCopy logical drive.

Repeat this step for each FlashCopy logical drive in the AIX array.

11. Re-create a new array. The **recreatevg** command reads the array data structure inside a logical drive and reconstructs it. The command allocates new physical logical drive identifiers (PIDs) to the FlashCopy logical drives and enables access to the FlashCopy logical drive for the selected host.

At the host prompt, type the following command and press Enter:

```
recreatevg -y logical drivegroupname -L /directoryname os_device_name
```

where:

- *logical drivegroupname* is the name that you want to assign to the FlashCopy array.
- *directoryname* is the name of the directory where you want to mount the FlashCopy logical drive.
- *os_device_name* is the operating system device name of the FlashCopy logical drive. If your AIX array contains more than one FlashCopy logical drive, add an *os_device_name* for each logical drive.

The array is re-created, and contains the FlashCopy logical drive or drives.

12. Mount the FlashCopy logical drive to its intended host. At the host prompt, type the following command and press Enter.

```
mount mount-point
```

where *mount-point* is the name of the filesystem that being mounted. Include the *directoryname* that was used in step 12.

13. Ensure that the logical drives are back online. At the host prompt, type the following command and press Enter.

```
df -k
```

A list of the mounted disks displays.

14. Use the FlashCopy logical drive with your backup application, for speculative testing, or with another application.
15. When the FlashCopy logical drive is no longer required, unmount the filesystem. At the host prompt, type the following command and press Enter.

```
umount mount-point
```

16. Delete the array, created in step 12, that contains the FlashCopy logical drives. At the host prompt, type the following commands and press Enter.

```
varyoffvg logical drivegroupname  
exportvg logical drivegroupname
```

where *logical drivegroupname* is the name of the FlashCopy array.

17. Disable the FlashCopy logical drive. Click **Logical Drive** → **FlashCopy** → **Disable** in the Subsystem Management window.

If you disable a FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and avoids any reduction in performance that might result by retaining the availability of an unused FlashCopy logical drive.

Reusing FlashCopy logical drives: Typically, after a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is needed.

Before you begin: Stop all I/O activity to the base logical drive, or suspend data transfer. This ensures that an accurate point-in-time image of the base logical drive is captured.

Perform the following steps to create a new point-in-time image of the same base logical drive:

1. Unmount the filesystems in the array on which the FlashCopy logical drive is to be based. At the host prompt, type the following command and press Enter.

```
umount mount-point
```

where *mount-point* is the name of the filesystem that is being unmounted.

2. Ensure that the host operating system recognizes the FlashCopy logical drive. At the host prompt, type the following command and press Enter.

```
cfgmgr
```

Several minutes might pass while the operating system accesses the drives. When the operation completes, a window opens that contains the following message:

```
Device nodes have been updated
```

The new logical drives are now available from the operating system.

3. Start the Storage Management software. The Enterprise Management window opens.

4. Launch a Subsystem Management window using one of the following methods:
 - Select the storage subsystem in either the Device Tree view or Device Table. Then, click the Manage Device toolbar button or click **Tools** → **Manage Device**.
 - Right-click the storage subsystem in the Device Tree view or Device Table, and select **Manage Device**.
 - Double-click a storage subsystem node in the Device Table.
 - Select the storage subsystem in the Device Tree view or the Device Table, and then press **Enter**.

The Subsystem Management window opens in a separate window.

5. In the Storage Management software, re-create the FlashCopy logical drive. Click **Logical Drive** → **FlashCopy** → **Re-create** in the Subsystem Management window.
6. Clear the PVID for the FlashCopy logical drives. At the host prompt, type the following command and press Enter.

```
chdev -l os_device_name -a pv=clear
```

where *os_device_name* is the operating system device name of the FlashCopy logical drive.

Repeat this step for each FlashCopy logical drive in the AIX array.

7. Re-create a new array. At the host prompt, type the following command and press Enter:

```
recreatevg -y logical_drivegroupname -L directoryname os_device_name
```

where:

- *logical_drivegroupname* is the name that you want to assign to the FlashCopy array.
- *directoryname* is the name of the directory where you want to mount the FlashCopy logical drive.
- *os_device_name* is the operating system device name of the FlashCopy logical drive. If your AIX array contains more than one FlashCopy logical drive, add an *os_device_name* for each logical drive.

The array is re-created and contains the FlashCopy logical drive or drives.

8. Mount the filesystem to its intended host. At the host prompt, type the following command and press Enter:

```
mount mount-point
```

where *mount-point* is the name of the filesystem that is being mounted. Include the *directoryname* that was used in step 7.

9. Ensure that the logical drives are back online. At the host prompt, type the following command and press Enter:

```
df -k
```

A list of the mounted disks is displayed.

10. Use the FlashCopy logical drive with your backup application or other application.

11. After the FlashCopy logical drive is no longer needed, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that might occur if the FlashCopy logical drive remains available.

Instructions for HP-UX: Logical Volume Manager

Use the following procedures to create FlashCopy logical drives on a host running HP-UX 11.0 (or higher) using Logical Volume Manager (LVM). Failure to complete the steps listed can result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives can either be created for repeated reuse (frequent or nightly backups) or for one-time use (speculative change or upgrade testing).

Creating a FlashCopy logical drive on HP-UX:

Before you begin: Read the information in “Using the script editor and command-line interface” on page 121.

Perform the following steps to create a FlashCopy logical drive:

1. Start the storage management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window, using one of the following methods:
 - Select the storage subsystem in the Device Tree view or from the Device Table. Then, click the **Manage Device** toolbar button, or click **Tools** → **Manage Device**.
 - Right-click the storage subsystem in the Device Tree view or from the Device Table and select **Manage Device**.
 - Double-click a storage subsystem node in the Device Table.
 - Select the storage subsystem in the Device Tree view or from the Device Table and press Enter.

The Subsystem Management window opens in a separate window.

Note: Stop the host application that is accessing the base logical drive and unmount the base logical drive. Unmounting the base logical drive does not apply when the base logical drive is the root disk of the host operating system.

3. When creating a FlashCopy logical drive based on a mounted file system, always run the **sync** command to flush the file system cache immediately before you create a FlashCopy logical drive.

At the host prompt, type the following command and press Enter.

```
sync
```

All unwritten file system buffers are flushed.

4. Within the Logical view of the Subsystem Management window, select a standard logical drive, and create a FlashCopy logical drive using the Create a FlashCopy Logical Drive wizard. This is accessed using a Subsystem Management window.
5. Assign a logical drive-to-LUN mapping between the FlashCopy logical drive and the host that will access the FlashCopy logical drive. Logical drive-to-LUN mappings can be defined using the Storage Partitioning wizard to quickly define a single storage partition. It guides you through the major steps to specify which host will access a logical drive, and the associated LUNs.
6. At the host prompt, type the following command and press Enter.

```
ioscan -fn
```

A list of the mapped devices that are recognized by the host is displayed.

Note: If the required device names are not displayed using this command, at the host prompt, type the following command and then press Enter.

```
insf
```

7. Remount the base logical drive to its original host.

Note: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive or reenable data transfer at this time.

8. Perform the following steps to import the FlashCopy logical drives into the Logical Volume Manager:
 - a. Create a new directory for the new array, as shown in the following example.

```
mkdir /dev/vg02
```

- b. Create a group node for the new array, as shown in the following example.

```
mknod /dev/vg02/group c 64 -0x020000
```

- c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter.

```
vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2
```

The following example addresses how to handle redundancy with multiple paths. The two paths or nodes represent the primary and alternate paths of the FlashCopy volume. Consider the following example:

```
vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
```

Note: You must verify that the `/dev/dsk` device files exist on the FlashCopy logical drive. Use the `SMdevices` utility or the `HP-UX ioscan` utility to perform the verification.

The system displays a warning indicating that a backup of the array being imported might not exist on the host. This message is only a warning and requires no response. The import continues and completes successfully.

The backup for this array is created later, when it is exported.

d. Activate the new array as shown in the following example:

```
vgchange -a y /dev/vg02
```

9. If a file system existed on the base logical drive, then it also exists on the FlashCopy logical drive. Before you mount the FlashCopy logical drive, however, perform a file system check to ensure that the file system is consistent, for example:

```
fscck /dev/vg02/1vol01
```

10. Mount the FlashCopy logical drive to its intended host.
11. Use the FlashCopy logical drive with your backup application, for speculative testing, or with another application.
12. Unmount the FlashCopy logical drive.
13. Disable the FlashCopy logical drive. Click **Logical Drive → FlashCopy → Disable** in the Subsystem Management window.

Disable the FlashCopy logical drive after it is no longer needed.

If you disable a FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and avoids any reduction in performance that might result by retaining the availability of an unused FlashCopy logical drive.

Reusing FlashCopy logical drives: Typically, after a FlashCopy logical drive is created it is disabled until a new point-in-time image of the same base logical drive is needed.

Before you begin: Stop all I/O activity to the base logical drive, or suspend data transfer. This ensures that an accurate point-in-time image of the base logical drive is captured.

Perform the following steps to create a new point-in-time image of the same base logical drive:

1. Start the storage management software. The Enterprise Management window opens.
2. Launch a Subsystem Management window using one of the following methods:
 - Select the storage subsystem in either the Device Tree view or from the Device Table. Then, click the **Manage Device** toolbar button or click **Tools→ Manage Device**.
 - Right-click the storage subsystem in the Device Tree view or from the Device Table, and select **Manage Device**.
 - Double-click a storage subsystem node in the Device Table.
 - Select the storage subsystem in the Device Tree view or from the Device Table, and then press Enter.

The Subsystem Management window opens in a separate window.

3. Unmount the base logical drive.
4. In the storage management software, re-create the FlashCopy logical drive. Click **Logical Drive** → **FlashCopy** → **Re-create** in the Subsystem Management window.
5. Remount the base logical drive to its original host.

Note: If you stopped I/O activity to the base logical drive or suspended a data transfer, resume I/O activity to the base logical drive or re-enable data transfer at this time.

6. Perform the following steps to import the FlashCopy logical drives into the Logical Volume Manager:
 - a. Create a new directory for the new array as shown in the following example.

```
mkdir /dev/vg02
```

- b. Create a group node for the new array as shown in the following example.

```
mknod /dev/vg02/group c 64 -0x020000
```

- c. Import the FlashCopy logical drive LUNs. At the host prompt, type the following command and press Enter.

```
vgimport /dev/vg02 FlashCopy-block-node-1 FlashCopy-block-node-2 ...
```

An alternate command can be used as shown in the following example.

```
vgimport /dev/vg02 /dev/dsk/c66t0d1 /dev/dsk/c69t0d1
```

Note: You must verify that the `/dev/dsk` device files exist on the FlashCopy logical drive. Use the `SMdevices` utility or the `HP-UX ioscan` utility to perform this verification.

A warning is displayed indicating that a backup of the array being imported might not exist on the host. This message is only a warning and can be ignored. The import continues and completes successfully.

The backup for this array is created when it is exported later.

- d. Activate the new array, as shown in the following example.

```
vgchange -a y /dev/vg02
```

7. If a file system existed on the base logical drive, then it also exists on the FlashCopy logical drive. However, before the FlashCopy logical drive can be mounted, submit it to a file system check.

Perform a file system check to ensure that the file system is consistent, for example:

```
fsck /dev/vg02/lvo101
```

8. Mount the FlashCopy logical drive to its intended host.

9. Use the FlashCopy logical drive with your backup application or other application.
10. Unmount the FlashCopy logical drive.
11. If the FlashCopy logical drive is no longer needed, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that might occur if the FlashCopy logical drive remains available.

Instructions for Solaris

Use the following procedures to create and reuse a FlashCopy logical drive using Solaris 2.6, 2.7 (Solaris 7), and 2.8 (Solaris 8). Failure to complete the steps listed can result in an inaccurate point-in-time image of the base logical drive.

FlashCopy logical drives can either be created for reuse (frequent or nightly backups) or for one-time use (speculative change or upgrade testing).

Creating a FlashCopy logical drive on Solaris:

Before you begin: Determine the following:

- If you plan to use the FlashCopy logical drive immediately, stop all I/O activity to the base logical drive. Also, unmount the base logical drives from the host to which they are currently mounted. This ensures that it will capture an accurate point-in-time image of the data on the base logical drive.
- If you plan to use the FlashCopy logical drive at a later date, do not unmount or stop I/O activity to the base logical drive at this time. You must perform these activities immediately before you use the FlashCopy logical drive for the first time.
- Read the information in “Using the script editor and command-line interface” on page 121.

Perform the following steps to create a FlashCopy logical drive:

1. Start the storage management software. The Enterprise Management window opens.
2. In the Logical view of the Subsystem Management window, select a standard logical drive and create a FlashCopy logical drive using the Create FlashCopy Logical Drive wizard. This is accessed using a Subsystem Management window.
3. If supported by the operating system, run the hot_add utility (or operating system-specific utility), or restart the host on which the FlashCopy will be used. This ensures that the host operating system recognizes the FlashCopy logical drive.

After you have created logical drives and defined logical drive-to-LUN mappings, run the hot_add utility to ensure that the operating system is aware of the newly-created logical drives, without having to restart the host. For information about which operating systems support the hot_add utility, see the Storage Manager software installation guide for your specific operating system, which is shipped on a CD with the product. You can also download it from one of the following Web pages:

ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt700
ssddom02.storage.ibm.com/techsup/webnav.nsf/support/fastt900

For the Solaris operating system, at the host prompt, type the following command and press Enter.

```
/etc/raid/bin/hot_add
```

After logical drives are created and logical drive-to-LUN mappings are defined, this step ensures that the operating system is aware of the newly-created logical drives, without having to restart the host.

4. Run the SMdevices utility to associate the LUN with a host operating system device and to ensure that the FlashCopy logical drive is recognized by the host. After you have created logical drives and defined logical drive-to-LUN mappings, run the SMdevices utility to ensure that the logical drive name and the operating system device name (assigned by the operating system) correlate.
5. If you plan to use the FlashCopy logical drive immediately, go to step 6. If you plan to use the FlashCopy logical drive at a later date, disable the FlashCopy logical drive now. Click **Logical Drive** → **FlashCopy** → **Disable** in the Subsystem Management window.

Note: If I/O activity to the base logical drive was stopped or data transfer was suspended, resume I/O activity to the base logical drive or re-enable data transfer at this time.

6. Mount the FlashCopy logical drive to its intended host.
7. Use the FlashCopy logical drive with your backup application, for speculative testing, or with another application.
8. Unmount the FlashCopy logical drive.
9. After the FlashCopy logical drive is no longer needed, disable or delete the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and avoids any reduction in performance that might occur if the FlashCopy logical drive remains available.

Reusing FlashCopy logical drives: Typically, after a FlashCopy logical drive is created, it is disabled until a new point-in-time image of the same base logical drive is needed.

Perform the following steps to create a new point-in-time image of the same base logical drive:

Note: Stop all I/O activity to the base logical drive at this point, or suspend data transfer. This ensures that an accurate point-in-time image of the base logical drive is captured.

1. Unmount the base logical drive.
2. In the storage management software, re-create the FlashCopy logical drive. Click **Logical Drive** → **FlashCopy** → **Re-create** from the menu in the Subsystem Management window.
3. Remount the base logical drive to its original host.
4. Mount the FlashCopy logical drive to its intended host.

Note: If you stopped I/O activity to the base logical drive or suspended data transfer, resume I/O activity to the base logical drive or re-enable data transfer at this time.

5. Use the FlashCopy logical drive with your backup application or with another application.
6. Unmount the FlashCopy logical drive.
7. When the FlashCopy logical drive is no longer needed, disable the FlashCopy logical drive.

If you disable the FlashCopy logical drive instead of deleting it, you can retain the FlashCopy logical drive and its associated FlashCopy repository logical drive. Then, when you need to create a different FlashCopy of the same base logical drive, you can re-create the disabled FlashCopy logical drive. This takes less time than creating a new FlashCopy logical drive and stops any reduction in performance that might occur if the FlashCopy logical drive remains available.

Using the script editor and command-line interface

Many storage management options that are available through the Subsystem Management window can be sent to the storage subsystem using statements in scripts. You can use the script editor to create or edit a script file, save a script file to the local disk of the Storage Management station, or load a script file from disk. You can also use the command-line interface (CLI) to issue individual commands to the scripting engine from the host operating system command shell or to call complete prewritten scripts.

For more information about how to use the command-line interface (CLI) and the script editor, click on the **Help** tab on the Enterprise management window and search for SMcli or script editor.

Checking FlashCopy status


Use this procedure to view a list of premium features on the storage subsystem and to verify that the FlashCopy feature has been enabled. It requires you to:

1. Check the premium options list.
2. View the FlashCopy icon to verify that the feature has been enabled.

Checking the premium options list

Perform the following steps to check the premium options list:

1. From the Subsystem Management window, click **Storage Subsystem** —> **Premium Features** —> **List**. Alternately, right-click and click **Premium Features** —> **List** (see Figure 32 on page 122).



SJ000827

Figure 32. Listing premium features

The List Premium Features window opens. It lists the following items:

- The Premium features that are enabled on the storage subsystem
- The Feature enable identifier

2. Verify that **FlashCopy Logical Drives:** indicates *Enabled* as shown in Figure 33.



SJ000828

Figure 33. Features list showing FlashCopy enabled

Figure 34 on page 123 shows an example where the FlashCopy copy feature is not enabled.



Figure 34. Feature list showing FlashCopy not enabled


See “Enabling FlashCopy” on page 94 for the procedure to enable FlashCopy.

3. Click **Close** to close the window.

Note: If you receive a Premium Features - Out of Compliance error message during a management session, use the Recovery Guru to resolve the problem.

Viewing the FlashCopy icon

Check the status of the FlashCopy feature icon at the bottom left of the device management GUI, as shown in Figure 35.



SJ000830

Figure 35. FlashCopy feature icon

The example in Figure 36 shows a disabled FlashCopy feature icon.



SJ000831


Figure 36. FlashCopy feature icon disabled

Mapping a FlashCopy drive to a host

Assign logical drive-to-LUN mappings between the FlashCopy logical drive and the host that will access the FlashCopy logical drive using the Mappings view of the Subsystem Management window. In some cases, depending on the host operating system and if any logical volume manager software is in use, mapping the same host to both a base logical drive and its associated FlashCopy logical drive can result in conflicts. For operating-system-specific instructions, see “Instructions for Solaris” on page 119 or “Instructions for HP-UX: Logical Volume Manager” on page 115.

Perform the following steps to map the FlashCopy logical drive to a host:


1. Open the Mappings view of the Management window. See Figure 37 on page 124. The newly-created FlashCopy logical drive displays in the undefined mapping section.



SJ000843

Figure 37. Undefined FlashCopy disk

2. Select the host or host group that you want to map the drive.
3. Right-click and select **Define Additional Mapping**, as shown in Figure 38.




SJ000844

Figure 38. Define Additional Mapping

4. The Define Additional Mapping window (Figure 39 on page 125) opens:
 - a. Select the FlashCopy drive.
 - b. Select the host or host group.
 - c. Set the LUN number.

d. Click **Add**.




SJ000845

Figure 39. Define Additional Mapping window

In the example shown in Figure 38 on page 124 and Figure 39, the FlashCopy logical disk DataVol-1 is made available to the backup server host group. This enables the backup server to access and mount the disk as part of its own file system. It also allows a local backup to be performed by the backup application software.

It is possible to map the FlashCopy logical disk to the same server that owns the base logical disk. However, immediately after creating the FlashCopy, the two logical disks appear exactly the same (a block-by-block copy). Many operating systems do not tolerate seeing an exact duplicate volume. You might need to complete other steps before you can access it. The mapping is shown in Figure 40 on page 126.

Note: If you use this FlashCopy on a regular basis (for example, for backup purposes), use the Disable FlashCopy and Re-create FlashCopy options to reuse the FlashCopy. Using these options preserves the existing mappings to the FlashCopy logical drive. For more information, see “Disabling a FlashCopy logical drive” on page 131 and “Recreating a FlashCopy logical drive” on page 133.



SJ000846

Figure 40. Mapped FlashCopy logical disk

Finally, use specific operating-system and host utilities to mount and use the mapped FlashCopy drive.

1. Run the host-based hot_add utility to register the FlashCopy logical drive with the host operating system.
2. Run the host-based SMdevices utility to associate the mapping between the physical device name and the logical drive name.
3. Mount the logical drive to the host.

For information about specific host operating-system procedures, see “Creating FlashCopy logical drives on AIX, HP-UX, and Solaris” on page 110.

Viewing the FlashCopy logical drive status

You can determine the status of the FlashCopy logical drive by viewing the icons that change depending of the state of the drive. The logical drive component property display is also useful in determining the state of the logical drives.

Use the FlashCopy Repository Logical Drive - Properties window to view the FlashCopy repository logical drive base and capacity properties. You can also use this window to specify the capacity percentage full and the action to be taken if the FlashCopy repository logical drive becomes full.

The progress of modification operations displays at the bottom of the window.

FlashCopy icon states

To view the FlashCopy icon, open the Storage Management Device Manager GUI Physical/Logical view. The icon states are shown in Figure 41 on page 127.

FlashCopy status representation in logical view		
Logical drive status icon	Mirror status	Icon
FlashCopy	Optimal	
	Disabled	
	Failed	
	Offline	
Repository	Optimal	
	Degraded	
	Failed	
	Full	
	Offline	
	Warning	
Free capacity	Spare capacity	


SJ000847

Figure 41. FlashCopy icon states

FlashCopy repository properties

Perform the following steps to view the FlashCopy repository properties:


1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** → **Properties**, or right-click and select **Properties**, as shown in Figure 42 on page 128.



SJ000848

Figure 42. Repository Logical Drive Properties

The FlashCopy Repository Logical Drive - Properties window opens, as shown in Figure 43 on page 129.




SJ000849

Figure 43. Base Repository Logical Drive Properties

The **Base** tab of the FlashCopy Repository Logical Drive - Properties window displays the following information for the selected FlashCopy repository logical drive:

- Logical drive name
 - Worldwide name
 - Status
 - Controller ownership
 - Capacity
 - RAID level
 - Modification priority
 - Associated base logical drive
 - Associated FlashCopy logical drive
3. Click the **Capacity** tab (Figure 44 on page 130) to view or set the following FlashCopy repository logical drive capacity properties:
- **FlashCopy repository logical drive percentage full**
Allows a threshold level to be set for the FlashCopy repository logical drive capacity. When the defined percentage is reached, a warning is issued indicating that the repository is nearing its capacity. The default percentage setting is 50% of the maximum capacity of the FlashCopy repository logical drive.
Use the **percent (%) full** field to define the percentage at which a warning is issued.



SJ000850

Figure 44. Repository drive warning threshold

- **Repository full policy**

When a FlashCopy repository logical drive reaches its capacity and becomes full, one of the following actions occurs:

- Fail FlashCopy logical drive
- Fail writes to base logical drive

If the FlashCopy repository logical drive is set to *Fail FlashCopy logical drive* (as shown in Figure 44) when it becomes full, its data is not recoverable and the FlashCopy cannot be accessed. In this case, the only available option is to delete the FlashCopy logical drive or re-create the FlashCopy logical drive to create a new point-in-time image.

If the FlashCopy repository logical drive is set to *Fail writes to the base logical drive*, the data is recoverable. However, the FlashCopy repository logical drive capacity must be increased so writes to the base logical drive are not rejected.

Attention: The *Fail writes to base logical drive* option of the repository full policy is not supported on AIX. Selecting this option might cause data loss on the base logical drive. If you are using AIX, select the *Fail FlashCopy logical drive* option, which is the default. Be sure to monitor the capacity of the FlashCopy repository logical drive, because you cannot access the FlashCopy if the repository logical drive becomes full.

For more information about increasing the storage capacity of an existing FlashCopy repository logical drive, see “Resizing a FlashCopy repository logical drive” on page 134.

Viewing the progress of a modification operation

The progress bar at the bottom of the FlashCopy Repository Logical Drive Properties window displays the progress of an operation. You can view the progress of the following operations:

- Copyback
- Reconstruction
- Initialization
- RAID level change
- Dynamic logical drive expansion
- Capacity increase
- Defragmentation
- Segment size change

Note: The storage management software cannot obtain progress information from the storage subsystem controllers if the network management connection to the controllers is down or if the storage subsystem is partially managed. For more information about a partially managed storage subsystem or an unresponsive controller or storage subsystem condition, see the Enterprise Management window online help.

Disabling a FlashCopy logical drive

If you no longer need a FlashCopy logical drive, you might want to disable it. As long as a FlashCopy logical drive is enabled, your storage subsystem performance can be impacted by copy-on-write activity that is directed to the associated FlashCopy repository logical drive. The copy-on-write activity concludes when you disable a FlashCopy logical drive.

If you disable a FlashCopy logical drive instead of deleting it, you can retain it and its associated repository for future use. When you need to create a different FlashCopy of the same base logical drive, you can use the Recreate option to reenable the previously disabled FlashCopy. This takes less time than creating a new FlashCopy.

When you disable a FlashCopy logical drive, note the following points:

- You cannot use that FlashCopy logical drive again until you use the Recreate option on that logical drive.
- Only that FlashCopy logical drive is disabled. All other FlashCopy logical drives remain functional.

If you do not intend to re-create a FlashCopy, you can delete that FlashCopy logical drive instead of disabling it.

Perform the following steps to disable a FlashCopy logical drive:

1. Select the FlashCopy logical drive. Right-click and select **Disable**, as shown in Figure 45 on page 132.




Figure 45. Choosing to disable the FlashCopy drive

2. The Disable FlashCopy Logical Drive confirmation window opens. See Figure 46. On this window, type Yes and click **OK** to begin the disable operation.




Figure 46. Disable FlashCopy confirmation window

The FlashCopy icon in the Physical/Logical view now displays as disabled, as shown in Figure 47 on page 133.




Figure 47. Icon showing the disabled FlashCopy logical drive

Recreating a FlashCopy logical drive

Recreating a FlashCopy logical drive takes less time than creating a new one. If you have a FlashCopy logical drive that you no longer need, instead of deleting it, you can reuse it (and its associated FlashCopy repository logical drive) to create a different FlashCopy logical drive of the same base logical drive.


When you re-create a FlashCopy logical drive, note the following points:

- The FlashCopy logical drive must be in either an optimal or a disabled state.
- All copy-on-write data on the FlashCopy repository logical drive is deleted.
- FlashCopy and FlashCopy repository logical drive parameters remain the same as the previously disabled FlashCopy logical drive and its associated FlashCopy repository logical drive. After you re-create a FlashCopy logical drive, you can change parameters on its associated FlashCopy repository logical drive through the appropriate menu options.
- The system retains the original names for the FlashCopy and FlashCopy repository logical drives. You can change these names, however, after the Recreate option completes.
- When using this option, the previously-configured FlashCopy name, parameters, and FlashCopy repository logical drive are used.

Perform the following steps to re-create a FlashCopy drive:

1. Select the FlashCopy logical drive. Right-click and select **Re-Create**, as shown in Figure 48 on page 134.
2. The Re-create FlashCopy Logical Drive window opens. Type Yes and click **OK**.

This command disables and re-creates the FlashCopy logical drive (if it had not been previously disabled) and displays it in the Logical view in an Optimal state. The creation timestamp shown on the FlashCopy Logical Drive Properties window is updated to reflect the new point-in-time image. Copy-on-write activity resumes to the associated FlashCopy repository logical drive.



SJ000854

Figure 48. Recreating a FlashCopy logical drive

Notes:

1. To use the Recreate option, the FlashCopy logical drive must be in either an optimal state or a disabled state.
2. If the FlashCopy logical drive is in an optimal state, the process first disables and then re-creates the FlashCopy logical drive. This process invalidates the current FlashCopy.

Resizing a FlashCopy repository logical drive

Use this option to increase the storage capacity of an existing FlashCopy repository logical drive. Typically, this option is used when a warning is received that the FlashCopy repository logical drive is in danger of becoming full.

You can achieve an increase in storage capacity by:

- Using free capacity available on the array of the FlashCopy repository logical drive.
- Adding unconfigured capacity (in the form of unused drives) to the array of the FlashCopy repository logical drive. Use this option when no free capacity exists on the array.

Note: A maximum of two drives can be added at one time to increase FlashCopy repository logical drive capacity.

The storage capacity of a FlashCopy repository logical drive cannot be increased if:

- 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.

Perform the following steps to resize a FlashCopy repository drive:

1. Select a FlashCopy repository logical drive in the Logical view of the Subsystem Management window.
2. Click **Logical Drive** → **Increase Capacity**, or right-click and select **Increase Capacity**.

Note: If no free capacity or unconfigured capacity is available, the Increase Capacity option is not available.

The Increase Repository Capacity window opens. You can see the FlashCopy repository logical drive name, the associated FlashCopy logical drive name, the associated base logical drive name, the current capacity, and the amount of free capacity that is available for the selected repository.

If free capacity is available, the maximum free space is shown in the **Increase Capacity by** field. If there is no free capacity on the array, the free space that is shown in the **Increase Capacity by** field is 0. Add drives to create free capacity on the array of the standard logical drive.

3. Use one of the following two methods to increase capacity:
 - Increase FlashCopy repository logical drive capacity using free capacity on the array of the FlashCopy repository logical drive:
 - a. Accept the final capacity increase or use the **Increase Capacity by** field to adjust the capacity. Click **OK**.
 - b. A confirmation window displays. Type Yes and click **OK** to continue.
The Logical view is updated. The FlashCopy repository logical drive with its capacity increased shows a status of Operation in Progress, together with its original capacity and the total capacity being added.
 - c. In addition, the Free Capacity node involved shows a reduction in capacity. If all of the free capacity is used to increase the logical drive size, then the Free Capacity node involved is removed from the Logical view.
 - Increase FlashCopy repository logical drive capacity by adding unconfigured capacity (drives) to the array of the FlashCopy repository logical drive:
 - a. If no unassigned drives are available and empty slots in the drive enclosures are available, insert new drives.
If no unassigned drives are available, and there are no empty slots available in the drive enclosures, install another drive enclosure and additional drives.
 - b. Select **Add Drives**.
The Increase Repository Capacity - Add Free Capacity window opens. Details about enclosure, slot, and usable capacity for the available free drives are displayed.

Note: The drives that are displayed have a capacity that is either equivalent to or larger than those that are already employed by the array.
 - c. Select a single drive, or two drives, to be added:
 - Press Ctrl and click **Enter** to select the nonadjacent drives.
 - Press Shift and click **Enter** to select the adjacent drives.
 - d. Click **Add**.

The Add Free Capacity window closes. Check the **Drives to add [enclosure, slot]** field to ensure that the correct drives are added.

- e. Accept the final capacity or use the **Increase Capacity by** field to adjust the capacity.
- f. Click **OK**.
- g. A confirmation window opens. Type Yes to confirm the operation, and click **OK** to continue.

The Logical view is updated. The FlashCopy repository logical drive having its capacity increased shows a status of *Operation in Progress*, together with its original capacity and the total capacity being added. In addition, the Free Capacity node involved in the increase shows a reduction in capacity.

If all of the free capacity is used to increase the logical drives size, then the Free Capacity node involved is removed from the Logical view.

If a Free Capacity node did not exist before you added capacity and not all of the capacity that is added will be used to increase the FlashCopy repository logical drives capacity, a new Free Capacity node is created and displayed in the Logical view.

Unassigned drives (unconfigured capacity) that are added to increase the FlashCopy repository logical drives capacity change in the Physical view to assigned drives, and become associated to the array of the FlashCopy repository logical drive.

4. View the progress of the capacity increase process. Select the FlashCopy repository logical drive. Click **Logical Drive → Properties**, or right-click and select **Properties**.

The FlashCopy Repository Logical Drive - Properties window opens. A progress bar at the bottom of the window indicates the status of the capacity increase.

Deleting a FlashCopy drive


Use this option to delete a FlashCopy logical drive that is no longer needed for backup or application testing purposes. This option results in an increase of free capacity in the array or additional unconfigured capacity.

Note:

- When you delete a logical drive, all data on the logical drive is lost. Back up the data and stop all I/O before you perform this operation, if necessary.
- If a file system is mounted on the logical drive, unmount it before you perform this operation.
- When you delete a base logical drive, the associated FlashCopy logical drive and FlashCopy repository logical drive are automatically deleted.

Perform the following steps to delete a FlashCopy drive:


1. Select the FlashCopy logical drive in the Logical view.
2. Click **Logical Drive → Delete**, or right-click and select **Delete**, as shown in Figure 49 on page 137.



SJ000855

Figure 49. Deleting the FlashCopy logical drive

3. The Delete FlashCopy Logical Drive window opens, as shown in Figure 50. Type Yes and click **OK**.



SJ000856

Figure 50. Delete FlashCopy Logical Drive Dialog window

The FlashCopy logical drive and FlashCopy repository logical drive are deleted, destroying any data that they might have contained.

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 drives that are associated with an array are removed, or when one or more drive enclosures lose power.

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

- The logical drive has an existing logical drive-to-LUN mapping, and drives that are associated with the logical drive are no longer accessible.
- The logical drive is participating in a Remote Volume Mirror as either a primary logical drive or a secondary logical drive, and drives that are associated with the logical drive are no longer accessible.
- The logical drive is a mirror repository logical drive, and drives that are associated with the logical drive are no longer accessible. The Recovery Guru has a special recovery procedure for this case. Two mirror repository logical drives are created together on the same array when the Remote Volume Mirroring premium 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 Volume 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 either a standard logical drive, base logical drive, FlashCopy repository logical drive, primary logical drive, secondary logical drive, or mirror repository logical drive.

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

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

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

Appendix A. AIX system requirements

This appendix lists the minimum hardware and software requirements that an AIX system must meet to be used in a FAST200, FAST500, FAST600, FAST700, or FAST900 storage subsystem that is managed by Storage Manager 8.3. A readme.txt file that contains the latest installation and user information about the storage management software, AIX filesets, and hardware components is located in the AIX FAST Installation CD. The most recent copy, which supercedes this document, is maintained on the following Web site:

www.ibm.com/pc/qtechinfo/MIGR-50177.html

After accepting the license agreement, proceed to the v8.3 link for the appropriate operating system.

Hardware requirements

The following hardware is required at a minimum.

- Any RS/6000® or pSeries server that supports FC 6227 or FC 6228 host bus adapters

Software requirements

The following software must be at the specified versions:

- AIX 4.3.3, AIX 5.1 or 5.2
- Filesets associated with IBM RDAC driver installations listed in Table 20, Table 21, or Table 22 on page 140.

RDAC installation requirements

The following filesets must be at the specified versions or later:

- For AIX 4.3.3:

Table 20. Filesets required for AIX 4.3.3 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	4.3.3.50
devices.fcp.disk.array.rte	4.3.3.84
devices.common.IBM.fc.rte	4.3.3.75
devices.pci.df1000f7.com	4.3.3.83
devices.pci.df1000f7.rte	4.3.3.75
devices.pci.df1000f9.rte	4.3.3.75
devices.scsi.sccarray.rte	4.3.3.50

- For AIX 5.1:

Table 21. Filesets required for AIX 5.1 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	5.1.0.0
devices.fcp.disk.array.rte	5.1.0.38
devices.common.IBM.fc.rte	5.1.0.10

Table 21. Filesets required for AIX 5.1 RDAC (continued)

PTF filesets	Version
devices.pci.df1000f7.com	5.1.0.35
devices.pci.df1000f7.rte	5.1.0.15
devices.pci.df1000f9.rte	5.1.0.15
devices.scsi.scarray.rte	5.1.0.0

- For AIX 5.2 with APAR IY35693:

Table 22. Filesets required for AIX 5.2 RDAC

PTF filesets	Version
devices.fcp.disk.array	5.2.0.2

Note: The AIX RDAC driver files are not included on the FAStT installation CD. Either install them from the AIX Operating Systems CD, if the correct version is included, or download them from the following Web site:

techsupport.services.ibm.com/server/fixes

For downloading instructions, see “Installing the RDAC driver” on page 20.

Appendix B. HP-UX system requirements

This appendix lists the minimum hardware and software requirements that an HP-UX system must meet to be used in a FAStT200, FAStT500, FAStT700, or FAStT900 storage subsystem that is managed by Storage Manager 8.3.

FAStT600:

HP-UX operating systems are not supported on the FAStT600 Fibre Channel Storage Server.

A readme.txt file that contains the latest installation and user information about the storage management software, and hardware components located in the HP-UX FAStT Installation CD. The most recent copy, which supercedes this document, is maintained on the following Web site:

www.ibm.com/pc/qtechinfo/MIGR-50177.html

After accepting the license agreement, proceed to the v8.3 link for the appropriate operating system.

The HP-UX system must be an HP 9000/Series 800 server with:

- 400 MHz processor or faster
- 0.7 MB available on /opt and root (or root-equivalent) privileges for installing SMutil and SMagent
- 47.5 MB available disk space on /opt and at least 95 MB available on /tmp for installation in an SMruntime environment

Ensure that the HP-UX host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.

- HP-UX 11.0 (64 bit) with the following patches:
 - PHKL_23939, version 1.0
 - PHKL_24004, version 1.0
 - PHKL_24027, version 1.0
 - PHKL_27364, version 1.0
 - PHKL_27003, version 1.0
 - To run HP-UX SDK for Java 2 Platform applications and applets using GUIs, you must also ensure that the HP C++ runtime libraries are installed on the HP-UX system; the latest version is available as patch PHSS_1658
- HP-UX 11.i (64 bit) with the following patches:
 - B.11.11.09
 - PHKL_23666, version 1.0
 - PHKL_26743, version 1.0
 - PHKL_27408, version 1.0
- For high-availability clusters of HP 9000/Series 800 computers, install the HP MC/Service Guard software package.

Ensure that the following maximum kernel parameters are configured, as shown in Table 23.

Table 23. HP-UX kernel parameter configuration requirements

Parameter	Description	Configuration
max_thread_proc 64	Maximum threads per process	1024
maxfiles	Soft-file limit per process	2048
maxuser	Influences other parameters	256 or greater
ncallout	Number of pending timeouts	4144

Appendix C. Solaris system requirements

This appendix lists the minimum hardware and software requirements that an Solaris system must meet to be used in a FAST200, FAST500, FAST600, FAST700, or FAST900 storage subsystem that is managed by Storage Manager 8.3. A readme.txt file that contains the latest installation and user information about the storage management software, and hardware components located in the Solaris FAST Installation CD. The most recent copy, which supercedes this document, is maintained on the following Web site:

www.ibm.com/pc/qtechenfo/MIGR-50177.html

After accepting the license agreement, proceed to the v8.3 link for the appropriate operating system.

The Solaris system must be an Sparc S20 processor with:

- 256 MB system memory
- CD-ROM drive
- Mouse or similar pointing device
- Ethernet network interface card
- 1 MB available on /opt and root (or root-equivalent) privileges for installing RDAC

Ensure that the Solaris host is running one of the following operating systems with the appropriate patches. Because patches can be superseded by more recent versions, refer to your operating system documentation or contact your operating system supplier to ensure that you have the correct patches.

- Solaris 2.6 with the following patches (minimum versions):
 - 105181-33
 - 105356-20
 - 106429-02
- Solaris 2.7 with the following patches (minimum versions):
 - 106541-23
 - 108376-42 (or later)
 - 107544-03
- 3. Solaris 2.8 requires the 06 Jumbo Patch. a. 108528-18 SunOS 5.8: kernel update patch b. 111293-04 SunOS 5.8: /usr/lib/libdevinfo.so.1 patch c. 111310-01 SunOS 5.8: /usr/lib/libdhcpagent.so.1 patch d. 111111-03 SunOS 5.8: /usr/bin/nawk patch e. 108987-12 SunOS 5.8: Patch for patchadd and patchrm
- Solaris 2.8 with the following patches (minimum versions):
 - 06 Jumbo Patch
 - 111293-04
 - 111310-01
 - 111111-03
 - 108987-12
- For high-availability clusters of Sparc S20 systems, install the Veritas Cluster Server software package. Check the Veritas Cluster Server (VCS) documentation for the latest patches.

Appendix D. MC/Service Guard configuration details

FAStT600:

MC/Service Guard configurations are not supported on the FAStT600 Fibre Channel Storage Server.

Note: The Hewlett Packard publication *Managing MC/ServiceGuard* suggests using **vgimport -m -s** with LVM commands during the configuration distribution process. It is important to consider, however, that the **vgimport** command does not preserve the primary and alternate paths of the imported volume groups.

When **vgimport** reconstructs the newly imported volume groups, entries are made in the `/etc/lvmtab` file. In this process of reconstruction, the system reorders disks in the file. The revised order of LUNs on the list causes LUNs to remain on non-preferred paths after failover. You might expect this condition to occur during distribution of the volume groups using the following LVM command:

```
vgimport -s -m /tmp/vg_group_name.map /dev/vg_group_name
```

This precaution applies only when you use the `-s` option to import devices or disks with redundant paths. The condition occurs because the `-s` option causes a search on the system for each disk when used with **vgimport**.

When the new list is prepared, the links that were previously designated as primary and alternate might not remain as they had been configured on node *A*.

Use the following procedure to correct the paths that were changed in this way when using **vgimport -m -s** with LVM commands.

Note: The following instructions are for a two-node cluster. You can easily expand or adapt these instructions to accommodate a cluster which has more than two nodes.

1. Identify and note the primary and alternate paths for each device by typing the following command:

```
# SMdevices
```

2. Compare the primary and alternate paths for each device on node *A* to Node *B* by using the **vgdisplay -v** command. As previously noted, an expected behavior of the **vgimport -m -s** command is that the primary and alternate paths change during export.

On node *A*:

```
#vgdisplay -v volume_group_name
```

On node *B*:

```
#vgdisplay -v volume_group_name
```

3. If the original primary path of a disk becomes an alternate path after the newly imported volume group entry is created in `/etc/lvmtab`, the order can be easily

reverted by using **vgreduce** to remove the wrong primary path and then using **vgextend** to add the path back again:

```
#vgreduce  
vg1 /dev/dsk/device_name#vgextend  
vg1 /dev/dsk/device_name
```

Appendix E. JNI host bus adapter settings

The following table detail settings for the various host bus adapter (HBA) cards for Sun Solaris. These cards are not plug-and-play with auto-configuration. Instead, you might need to change the settings or bindings.

Table 24. Configuration file name: /kernel/drv/fca-pci.conf

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	
ip_disable = 0;	ip_disable = 1;	
failover = 0;	failover = 30;	
busy_retry_delay = 5000;	busy_retry_delay = 5000;	
link_recovery_delay = 1000;	link_recovery_delay = 1000;	
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;	
def_hba_binding = "fca-pci*";	def_hba_binding = "nonjni";	
def_wwnn_binding = "\$xxxxxx"	def_wwnn_binding = "xxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxx"	Same as the original entry.	Uncomment the line.
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"	Uncomment the line.
Will be added by reconfigure script	target0_hba="fca-pci0" target0_wwpn="controller_wwpn";	
Will be added by reconfigure script	name="fca-pci" parent="physical path" unit-address="#"	
Will be added by reconfigure script	target0_hba="fca-pci1" target0_wwpn="controller_wwpn";	

Note: You might need to run the /etc/raid/bin/genjnicnf reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicnf
```

Table 25. Configuration file name: /kernel/drv/jnic146x.conf

Original value	New value	Comments
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 30;	Uncomment the line.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	Same as the original entry.	Uncomment the line.

Table 25. Configuration file name: /kernel/drv/jnic146x.conf (continued)

Original value	New value	Comments
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "<controller_wwpn>"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "<controller_wwpn>"	

Note: You might need to run the /etc/raid/bin/genjnicconf reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicconf
```

Table 26. Configuration file name: /kernel/drv/jnic.conf

Original value	New value	Comment
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 30;	Uncomment the line.
JniCreationDelay = 5;	JniCreationDelay = 10;	Uncomment the line.
def_wwnn_binding = "\$xxxxxxxxxxx"	def_wwnn_binding = "xxxxxxxxxxx"	Uncomment the line.
def_wwpn_binding = "\$xxxxxxxxxxx"	Same as the original entry.	Uncomment the line.
Add	target0_hba = "jnic146x0";	
Add	target0_wwpn = "controller_wwpn"	
Add	target1_hba = "jnic146x1";	
Add	target1_wwpn = "controller_wwpn"	

Note:

- You might need to run the /etc/raid/bin/genjnicconf reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicconf
```

- Set *portEnabled* = 1; only when you see JNI cards entering non-participating mode in the /var/adm/messages file. Under that condition:

- Set *FcPortCfgEnabled* = 1;
- Restart the host.
- Set *FcPortCfgEnabled* = 0;
- Restart the host again.

When you have done so, check /var/adm/messages to be sure that it sets the JNI cards to Fabric or Loop mode.

Table 27. Configuration file name: /kernel/drv/fcaw.conf

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	

Table 27. Configuration file name: /kernel/drv/fcaw.conf (continued)

Original value	New value	Comments
ip_disable = 0;	ip_disable=1;	
failover = 0;	failover =30;	
busy_retry_delay = 5000;	busy_retry_delay = 5000;	
link_recovery_delay = 1000;	link_recovery_delay = 1000;	
scsi_probe_delay = 5000;	scsi_probe_delay = 5000;	
def_hba_binding = "fcaw*";		Uncomment the line.
def_wwnn_binding = "\$xxxxxx"	def_wwnn_bindindef_hba_binding = "nonjni"; g = "xxxxxx"	Uncomment the line.
def_wwnn_binding = "\$xxxxxx"	Same as the original entry.	Uncomment the line.
Will be added by reconfigure script	name="fcaw" parent="<physical path>"unit-address="<#>"	
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn="<controller wwpn>";	
Will be added by reconfigure script	name="fcaw" parent="<physical path>"unit-address="<#>"	
Will be added by reconfigure script	target0_hba="fcaw0" target0_wwpn="<controller wwpn>";	

Note: You might need to run the /etc/raid/bin/gen SCSIconf reconfigure script from the shell prompt.

```
# /etc/raid/bin/gen SCSIconf
```

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "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 publications. IBM may make improvements or changes (or both) in the product(s) or program(s) (or both), described in this publication at any time without notice.

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

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

Trademarks

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

IBM
AIX
e(logo)
FlashCopy

Netfinity
RS/6000
TotalStorage

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

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

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

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

Glossary

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

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

This glossary also includes terms and definitions from:

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

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

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

See also

Refers you to a related term.

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

accelerated graphics port (AGP). A bus specification that gives low-cost 3D graphics cards faster access to main memory on personal computers than the usual 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 (I/Os) between the internal bus of the host system and the external Fibre Channel 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 an SNMP-TCP/IP network-managing environment.

AGP. See *accelerated graphics port*.

AL_PA. See *arbitrated loop physical address*.

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

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

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

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

AWT. See *Abstract Windowing Toolkit*.

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

BIOS. See *basic input/output system*.

BOOTP. See *bootstrap protocol*.

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

bridge. A SAN device that provides physical and transport conversion, such as Fibre Channel to SCSI bridge.

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

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

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

CDPD. See *cellular digital packet data*.

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

CGA. See *color graphics adapter*.

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

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

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

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

CRC. See *cyclic redundancy check*.

CRT. See *cathode ray tube*.

cyclic redundancy check (CRC). (1) 1) A redundancy check in which the check key is generated by a cyclic algorithm. (2) 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*.

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

DHCP. See *Dynamic Host Configuration Protocol*.

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

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

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

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 N_Port Identifier for the FC device. It is not used in the FC-SCSI hardware path ID. It is required to be the same for all SCSI targets logically connected to an FC adapter.

DRAM. See *dynamic random access memory*.

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

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

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

ECC. See *error correction coding*.

EEPROM. See *Electrically Erasable Programmable Read-Only Memory*.

EGA. See *enhanced graphics adapter*.

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

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

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

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

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

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

ESD. See *electrostatic discharge*.

ESDI. See *enhanced small disk interface*.

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

F_Port. A port that supports an N_Port on a Fibre Channel switch.

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

Fibre Channel. A bi-directional, full-duplex, point-to-point, serial data channel structured for high performance capability. Physically, Fibre Channel interconnects devices, such as host systems and servers, FC hubs and disk arrays, through ports, called N_Ports, in one of three topologies: a point-to-point link, an arbitrated loop, or a cross point switched network,

which is called a fabric. FC can interconnect two devices in a point-to-point topology, from two to 126 devices in an arbitrated loop. FC is a generalized transport mechanism that can transport any existing protocol, such as SCSI, in FC frames.

Fibre Channel Protocol for SCSI (FCP). A high-level Fibre Channel mapping layer (FC-4) that uses lower-level Fibre Channel (FC-PH) services to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target across the FC link by using FC frame and sequence formats.

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

FRU. See *field replaceable unit*.

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

GPIB. See *general purpose interface bus*.

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

GUI. See *graphical user interface*.

HBA. See *host bus adapter*.

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

host. 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.

host bus adapter (HBA). An interface between the Fibre Channel network and a workstation or server.

host computer. See *host*.

host group. The collection of HBAs and NASs in a fabric discovered by the SANavigator tool and displayed with a yellow background on the Physical and Data Path Maps.

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.

IC. See *integrated circuit*.

IDE. See *integrated drive electronics*.

In-band. Transmission of management protocol over the Fibre Channel transport.

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

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

integrated circuit (IC). Also known as a *chip*. 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.

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

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

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

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

IP address. See *Internet Protocol address*.

IPL. See *initial program Load*.

IRQ. See *interrupt request*.

ISA. See *Industry Standard Architecture*.

ISDN. See *Integrated Services Digital Network*.

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

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

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

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

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

loop address. The unique ID of a node in Fibre Channel loop topology sometimes referred to as a Loop ID.

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

loop port (FL_Port). An N-Port or F-Port that supports arbitrated loop functions associated with an arbitrated loop topology.

LUN. See *logical unit number*.

MAC. See *medium access control*.

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

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

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

MCA. See *micro channel architecture*.

MIB. See *management information base*.

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

MIDI. See *musical instrument digital interface*.

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

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

NDIS. See *network device interface specification*.

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

network management station (NMS). In the Simple Network Management Protocol (SNMP), a station that executes management application programs that monitor and control network elements.

NMI. See *non-maskable interrupt*.

NMS. See *network management station*.

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

N_Port. A node 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.

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

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

NVS. See *nonvolatile storage*.

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

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

ODM. See *Object Data Manager*.

out-of-band. Transmission of management protocols outside of the Fibre Channel network, typically over Ethernet.

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

PDF. See *portable document format*.

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

performance events. Events related to thresholds set on SAN performance.

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

port. The hardware entity that connects a device to a Fibre Channel topology. A device can contain one or more ports.

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

private loop. A freestanding Arbitrated Loop with no fabric attachment.

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

PTF. See *program temporary fix*.

RAM. See *random-access memory*.

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

RDAC. See *redundant disk array controller*.

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

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

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

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

RGB. See *red, green, blue*.

ROM. See *read-only memory*.

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

RVSD. See *recoverable virtual shared disk*.

SAN. See *storage area network*.

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

SCSI. See *small computer system interface*.

segmented loop ports (SL_Ports). SL_Ports allow you to divide a Fibre Channel Private Loop into multiple segments. Each segment can pass frames around as an independent loop and can connect through the fabric to other segments of the same loop.

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

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 ports*.

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

SNMP. See *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*.

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

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

SRAM. See *static random access memory*.

SSA. See *serial storage architecture*.

static random access memory (SRAM). Random access memory based on the logic circuit known as flip-flop. It is called *static* because it retains a value as long as power is supplied, unlike dynamic random access memory (DRAM), which must be regularly refreshed. It is however, still volatile, meaning that it can lose its contents when the power is switched off.

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

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

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

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

SVGA. See *super video graphics array*.

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

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

switch group. A switch and the collection of devices connected to it that are not in other groups. Switch Groups are discovered by the SANavigator tool and displayed with a gray background on the Physical and Data Path Maps.

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

TCP. See *Transmission Control Protocol*.

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

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

TFT. See *thin-film transistor*.

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

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

TL_Ports. See *translated loop port*.

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

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

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

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

trap recipient. Receiver of a forwarded SNMP trap. Specifically, a trap receiver is defined by an 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*.

user action events. Actions that the user takes, such as changes in the SAN, changed settings, and so on. Each such action is considered a User Action Event.

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

VGA. See *video graphics adapter*.

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

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

VRAM. See *video random access memory*.

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

Worldwide Name (WWN). A registered, unique 64-bit identifier assigned to nodes and ports.

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

WWN. See *worldwide name*.

XGA. See *eXtended graphics array*.

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

Index

A

- about this document xi
- access volumes 7
- address
 - for e-mail comments xxi
 - for reader comment form xxi
 - of IBM xxi
 - of the IBM director of licensing 151
- agent software
 - package 4
- AIX
 - client software 15
 - configuration of storage subsystems 16, 21
 - direct-attached configuration 19
 - firmware
 - requirements 13
 - upgrading 17
 - FlashCopy
 - creating a FlashCopy logical drive 110
 - reusing a FlashCopy logical drive 113
 - hardware requirements 139
 - identifying
 - controller ID numbers 29
 - device names and bus numbers 30
 - FlashCopy volume types 35
 - logical drives by operating system device names 33
 - NVSRAM
 - requirements 13
 - upgrading 17
 - operating system requirements 8, 13
 - RDAC driver
 - configuring 21
 - installing 20, 21
 - requirements 20, 139
 - verifying 20
 - requirements
 - client software 15
 - firmware 13
 - hardware 13, 139
 - operating system 8
 - RDAC driver 20, 139
 - software 139
 - system 139
 - restrictions 5, 7, 14, 97, 105, 130
 - SAN-attached configuration 19
 - SMclient 16
 - SMruntime 16
 - software requirements 139
 - system requirements 139
 - volumes, redistributing in case of failure 39
- alert notifications, configuring 77
- attributes
 - dac 23, 27
 - dar 22, 26
 - definitions 22
 - hdisk 24, 27, 28

- attributes (*continued*)
 - LUN 24, 27, 28
 - RDAC driver 21
- audience of this document xi
- Auto Volume Transfer (AVT) 53, 72

B

- bus numbers and device names, identifying 30, 52

C

- capture state information feature 79
- catastrophic failure 79
- changing the host type 73
- client software
 - AIX 15, 16
 - HP-UX 42, 43
 - installing 42
 - package 3, 4
 - Solaris 57, 58
 - uninstalling 61
- cluster services
 - AIX requirements 83, 84
 - configurations 74
 - HACMP ES and ESCRM 81
 - hardware requirements 82
 - HP-UX requirements 81
 - MC/Service Guard 81
 - PSSP with GPFS 81
 - Solaris requirements 81, 83
 - system dependencies 81
 - Veritas Cluster Server 81
 - Veritas Volume Manager 81, 82
- commands
 - fget_config 32
 - lsattr 33
 - lsdev 31
- comments about this document, how to send xxi
- configuration details, MC Service Guard 145
- configuration recovery 79
- configuring
 - alert notifications 77
 - direct-attached 19, 48, 64
 - RDAC driver 21, 82
 - SAN-attached 19, 48, 64
 - Solaris with JNI host bus adapter cards 65, 66
 - storage subsystems
 - AIX 16, 19
 - HP-UX 46, 48
 - procedures 74
 - Solaris 61, 64
- controller state capture 79
- controllers
 - firmware 1, 20, 46, 62
 - ID numbers, identifying 29, 51
 - IP addresses 8

controllers (*continued*)
properties 51

D

dac attributes 23, 27
dar attributes 22, 26
DCE (dynamic capacity expansion) 36
default partitioning, Solaris devices 63
devices
 identification 28
 names and bus numbers, identifying 30, 52
 Solaris, default partitioning for 63
direct (out-of-band) management method 7
direct-attached configuration 19, 48, 64
disk configuration information 79
documentation
 documents 1
 FASTT xii
 FASTT Storage Manager Version 8.3 xiii
 FASTT200 xviii
 FASTT500 xvii
 FASTT600 xvi
 FASTT700 xv
 FASTT900 xiv
 files 1
 online help xx, 1, 76, 78
 related xix
 Web sites xxi, 1
DVE (dynamic volume expansion) 36
dynamic capacity expansion (DCE) 36
dynamic volume expansion (DVE) 36

E

e-mail address xxi, 77
edition notice ii
enabling multipath I/O with PV-links 49
Enterprise Management window
 description 4
 online help 1
 starting 74

F

failure support
 cluster services 81
 RDAC driver 5
 redistributing volumes 39, 53, 72
FASTT documentation xii
FASTT Storage Manager
 related documents xix
 user tasks xiii
 Version 8.3 library xiii
FASTT Storage Manager Field Tool 79
FASTT, introduction to 8
FASTT200 Storage Server library xviii
FASTT500 Storage Server library xvii
FASTT600
 FC Storage Server library xvi

FASTT600 (*continued*)
 restrictions 3, 8, 19, 20, 41, 64, 65, 81, 83, 84, 93,
 141, 145

FASTT700 Storage Server library xv

FASTT900 FC Storage Server library xiv

fget_config command 32

filesets 139

firmware

 AIX

 requirements 13

 upgrading 17

 verifying 20

 description 1

 HP-UX

 requirements 41

 upgrading 47

 Solaris

 requirements 55

 upgrading 62

 versions 46, 62

FlashCopy 93

 AIX

 creating a FlashCopy logical drive 110

 reusing a FlashCopy logical drive 113

 checking status 121

 command line interface 121

 enabling 94

 HP-UX

 creating a FlashCopy logical drive 115

 reusing a FlashCopy logical drive 117

 logical drive

 creating 100

 deleting 136

 disabling 131

 mapping to a host 123

 re-creating 133

 viewing status 126

 overview 93

 planning a logical drive 95

 recovering a missing logical drive 138

 repository logical drive

 calculating expected overhead 98

 estimating capacity 98

 estimating repository life 99

 properties 127

 re-sizing 134

 repository full policy 97, 130

 setting options 95

 viewing failure settings 95

 script editor 121

 script scenarios 93

 Solaris

 creating a FlashCopy logical drive 119

 reusing a FlashCopy logical drive 120

 viewing progress of a modification operation 131

H

HACMP ES and ESCRM 81

hardware installation process xi

- hardware requirements
 - AIX 13, 139
 - HP-UX 41, 141
 - Solaris 55, 143
- hdisk attributes 24, 27, 28
- help xx, 1, 76, 78
- heterogeneous environment 74
- high-availability cluster services 81
- host software
 - HP-UX 44, 45
 - packages 3
 - Solaris 59, 61
- host type 73
- host-agent (in-band) management method 5
- how this document is organized xx
- how to send your comments xxi
- HP-UX
 - client software 42
 - cluster services, requirements 81
 - configuration of storage subsystems 46
 - direct-attached configuration 48
 - firmware
 - requirements 41
 - upgrading 47
 - FlashCopy
 - creating a FlashCopy logical drive 115
 - reusing a FlashCopy logical drive 117
 - hardware requirements 41, 141
 - host software
 - installing 44
 - requirements 44
 - identifying
 - controller ID numbers 51
 - device names and bus numbers 52
 - NVSRAM
 - requirements 41
 - upgrading 47
 - operating system requirements 8
 - requirements
 - client software 42
 - cluster services 81
 - firmware 41
 - hardware 41, 141
 - host software 44
 - NVSRAM 41
 - operating system 8
 - software 141
 - system 141
 - SAN-attached configuration 48
 - SMagent 44
 - SMclient 43
 - SMruntime 43
 - SMutil 45
 - software requirements 141
 - system requirements 141
 - volumes, redistributing in case of failure 53
- HyperTerminal version 9

I

- IBM
 - address xxi
 - director of licensing address 151
 - e-mail address xxi, 77
- identifying
 - AIX
 - controller ID numbers 29
 - device names and bus numbers 30
 - logical drives by operating system device names 33
 - HP-UX
 - controller ID numbers 51
 - device names and bus numbers 52
- in-band (host-agent) management method 5
- installation process xi
- installing
 - AIX
 - client software 15, 16
 - RDAC driver 20
 - client software
 - AIX 15
 - HP-UX 42
 - SMclient 16, 43, 58
 - SMruntime 16, 43, 57
 - Solaris 57
 - host software
 - HP-UX 44
 - SMagent 44, 59
 - SMutil 45, 60
 - Solaris 59
 - HP-UX
 - client software 42, 43
 - host software 44
 - SMutil 45
 - post-installation tasks 73
 - sequence of 3
 - Solaris
 - client software 57, 58
 - host software 59, 60
 - RDAC driver 69
- IP addresses for FAS*T* controllers 8

J

- JNI
 - cards 65, 66
 - settings 147

L

- lsattr command 33
- lsdev command 31
- LUN attributes 27
- LUNs
 - and access volumes 7
 - attributes 24, 27, 28

M

- management method
 - direct (out-of-band) 7
 - host-agent (in-band) 5
- management station 1
 - See storage management station
- MC Service Guard 145
- MC/Service Guard 81
- missing FlashCopy logical drives 138
- multipath I/O with PV-links 49

N

- new features 2
- notice, edition ii
- notices
 - general 151
- NVSRAM
 - AIX
 - requirements 13, 14
 - HP-UX
 - requirements 41
 - upgrading 47
 - Solaris
 - requirements 55
 - upgrading 62

O

- online help xx, 1, 76, 78
- operating system
 - AIX 8, 13, 139
 - HP-UX 8, 141
 - requirements 41
 - requirements
 - AIX 13
 - HP-UX 41
 - Solaris 55
 - Solaris 8, 143
 - requirements 55
- organization of this document xx
- out-of-band (direct) management method 7
- overview of heterogeneous hosts 74

P

- partitioning for Solaris devices (default) 63
- premium features 2
 - FlashCopy 93
 - Remote Mirror Option 93
 - Storage Partitioning 10
- prerequisites
 - AIX
 - client software 15
 - firmware 13
 - hardware 13, 139
 - operating system 8
 - RDAC driver 20, 139
 - software 139
 - system 139

prerequisites (continued)

- cluster services 82
- HP-UX
 - client software 42
 - cluster services 81
 - firmware 41
 - hardware 41, 141
 - host software 44
 - NVSRAM 41
 - operating system 8
 - software 141
 - system 141
- Solaris
 - client software 57
 - cluster services 81, 83
 - firmware 55
 - hardware 55, 143
 - host software 59
 - NVSRAM 55
 - operating system 8
 - RDAC driver 59
 - software 143
 - system 143
- products, developed 151
- PSSP with GPFS 81
- PV-links 49

R

- RDAC driver
 - AIX
 - configuring 21
 - installing 20, 21
 - requirements 20, 139
 - attributes 21
 - description 5
 - IDs 82
 - installing 69
 - Solaris
 - installing 69
 - requirements 59
 - uninstalling 61
- reader comment form xxi
- readme.txt file 1
- recover configuration 79
- redistributing volumes in case of failure 39, 53, 72
- Remote Mirror Option 93
- renaming storage subsystem 78
- requirements
 - AIX
 - client software 15
 - firmware 13
 - hardware 13, 139
 - operating system 8
 - RDAC driver 20, 139
 - software 139
 - system 139
 - cluster services 82
 - HP-UX
 - client software 42
 - cluster services 81

- requirements *(continued)*
 - HP-UX *(continued)*
 - firmware 41
 - hardware 41, 141
 - host software 44
 - NVSRAM 41
 - operating system 8
 - software 141
 - system 141
 - Solaris
 - client software 57
 - cluster services 81, 83
 - firmware 55
 - hardware 55, 143
 - host software 59
 - NVSRAM 55
 - operating system 8
 - RDAC driver 59
 - software 143
 - system 143
- resources
 - documents 1
 - e-mail address xxi, 77
 - files 1
 - online help xx, 1, 76, 78
 - Web sites xxi, 1
- restrictions
 - AIX 5, 7, 14
 - FAST600 3, 8, 19, 20, 41, 64, 65, 81, 83, 84, 93, 141, 145
 - in-band management 5

S

- SAN-attached configuration 19, 48, 64
- sending your comments to IBM xxi
- services offered in the U.S.A. 151
- setting
 - IP addresses 8
 - NMS for SNMP notification 77
- SMagent
 - description 4
 - HP-UX installing 44
 - installing 59
 - Solaris installing 59
- SMclient
 - AIX installing 16
 - description 4
 - HP-UX installing 43
 - Solaris installing 58
- SMruntime
 - AIX installing 16
 - HP-UX installing 43
 - Solaris installing 57
- SMutil
 - description 5
 - HP-UX installing 45
 - Solaris installing 60
- SNMP notification 77
- software installation process xi

- software package
 - RDAC 5
 - SMagent 4
 - SMclient 4
 - SMutil 5
- software requirements
 - AIX 139
 - HP-UX 141
 - Solaris 143
- Solaris
 - client software 57, 61
 - cluster services requirements 81, 83
 - configuration of storage subsystem 61
 - devices, default partitioning 63
 - direct-attached configuration 64
 - firmware
 - requirements 55
 - upgrading 62
 - FlashCopy
 - creating a FlashCopy logical drive 119
 - reusing a FlashCopy logical drive 120
 - hardware requirements 55, 143
 - host software 59, 61
 - JNI host bus adapter cards 65, 66
 - NVSRAM
 - requirements 55
 - upgrading 62
 - operating system requirements 8
 - RDAC driver
 - installing 69
 - requirements 59
 - uninstalling 61
 - requirements
 - client software 57
 - cluster services 81
 - firmware 55
 - hardware 55, 143
 - host software 59
 - NVSRAM 55
 - operating system 8
 - RDAC driver 59
 - software 143
 - system 143
 - SAN-attached configuration 64
 - SMagent
 - installing 59
 - SMclient
 - installing 58
 - SMruntime
 - installing 57
 - SMutil
 - installing 60
 - software requirements 143
 - system requirements 143
 - volumes, redistributing in case of failure 72
 - starting the Enterprise Management window 74
 - starting the Subsystem Management window 77
 - storage management station 1
 - Storage Manager software
 - installation sequence 3
 - introduction 1

- Storage Manager software *(continued)*
 - software packages 3
 - uninstalling, Solaris 61
- Storage Partitioning premium feature 10
- storage subsystem
 - cluster services 81
 - configuring
 - AIX 16
 - HP-UX 46
 - procedures 74
 - Solaris 61
 - introduction 1
 - logical drives
 - identifying by operating system device names 33
 - management methods
 - direct (out-of-band) 7
 - host-agent (in-band) 5
 - management tasks 78
 - renaming 78
- storage subsystem profile, saving 79
- Subsystem Management window
 - description 4
 - figure of 77
 - online help 2
 - starting 77
- switch
 - in a SAN-attached configuration 19, 48, 64
 - zoning 19, 48, 64
- system requirements
 - AIX 139
 - HP-UX 141
 - Solaris 143

T

- tasks by documentation title xii
- trademarks 151

U

- uninstalling Storage Manager software, Solaris 61
- upgrading
 - AIX
 - NVSRAM 17
 - HP-UX
 - firmware 47
 - NVSRAM 47
 - Solaris
 - firmware 62
 - NVSRAM 62
- user tasks
 - FAStT Storage Manager xiii
 - FAStT200 Fibre Channel Storage Server xviii
 - FAStT500 Fibre Channel Storage Server xvii
 - FAStT600 Fibre Channel Storage Server xvi
 - FAStT700 Fibre Channel Storage Server xv
 - FAStT900 Fibre Channel Storage Server xiv
 - related documents xix
- utility software
 - package 5

V

- variable settings for JNI adapter cards 66
- Veritas Cluster Server 81
- Veritas Volume Manager 81, 82
- volumes
 - redistributing in case of failure 39, 53, 72

W

- Web sites xxi, 1, 14, 20, 42, 83
- who should
 - read this document xi
 - use this document xi

Z

- zoning switches 19, 48, 64

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage FAStT Storage Manager Version 8.3
Installation and Support Guide for AIX, HP-UX, and Solaris

Publication No. GC26-7521-01

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<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"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<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
RCF Processing Department
Dept. M86/Bldg. 050-3
5600 Cottle Road
San Jose, CA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

GC26-7521-01



Spine information:



IBM TotalStorage FAS*StT* Storage
Manager Version 8.3

Installation and Support Guide for AIX, HP-UX, and
Solaris