

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 115.

Fourth Edition (December 2003)

This edition replaces GC26-7521-02.

© **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
FAST600 Fibre Channel Storage Server library.	xiv
FAST900 Fibre Channel Storage Server library	xv
FAST700 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	xxii
Chapter 1. Introduction	1
Overview of Storage Manager 8.3	1
The storage management station	1
Storage Manager 8.3 online help.	2
Changes in this edition	2
Storage Manager 8.3 premium features	2
FAST600 restrictions	4
Overview of Storage Manager 8.3 software packages	5
SMruntime software package	5
SMclient software package	5
SMagent software package.	6
SMutil software package.	6
RDAC	6
Software installation sequence	6
Storage subsystem management methods	7
Host-agent (in-band) management method	7
Direct (out-of-band) management method	9
Installation process diagram	11
Operating system requirements.	13
Setting up IP addresses for FAST storage controllers	14
Chapter 2. Installing storage management station software on AIX systems 17	
Hardware and firmware requirements	17
AIX restrictions	18
Installing the client software	19
Prerequisites.	19
Installing SMruntime	20
Installing SMclient.	20
Performing the initial configuration of storage subsystems on AIX hosts	20
Upgrading FAST firmware and NVSRAM	21
Setting up an AIX host group.	22
Creating direct and SAN-attached configurations	23
Creating a direct-attached configuration	23

Creating a SAN-attached configuration	24
Installing the RDAC driver	24
AIX configuration information	25
Viewing and setting attributes of the RDAC driver for AIX	25
Viewing Object Data Manager attributes in AIX	30
Changing ODM attribute settings in AIX	31
Verifying the installation and configuration of AIX hosts	33
Identifying the controller ID numbers	33
Identifying device names and bus numbers	34
Identifying logical drives by operating system device names	40
Identifying FlashCopy volume types	42
Using fast I/O failure for fibre-channel devices	43
Using dynamic tracking of fibre-channel devices	44
Using dynamic capacity expansion and dynamic volume expansion	44
Performing a dynamic capacity expansion operation	45
Performing a dynamic volume expansion operation	45
Resolving disk array errors	45
Redistributing volumes in case of failure	47

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

Hardware and firmware requirements	49
Installing the client software	50
Prerequisites	51
Installing SMruntime	51
Installing the SMclient	51
Installing the host software	52
Prerequisites	52
Installing the SMagent (optional)	52
Installing SMutil (required)	53
Performing the initial configuration of storage subsystems on HP-UX hosts	54
Adding storage subsystems to SMclient	54
Updating NVSRAM and FASTT firmware	55
Setting up an HP-UX host group	55
Creating direct-attached and SAN-attached configurations	56
Creating a direct-attached configuration	57
Creating a SAN-attached configuration	57
Enabling multipath I/O with PV-links	57
Using PV-links: Method 1	58
Using PV-links: Method 2	60
Redistributing volumes in case of failure	63

Chapter 4. Installing storage management station software on Solaris systems

Hardware and firmware requirements	65
Installing the client software	67
Prerequisites	67
Installing SMruntime	67
Installing SMclient	68
Installing host software on Solaris hosts	69
Prerequisites	69
Installing SMagent (optional)	70
Installing SMutil (required)	70
Uninstalling Storage Manager 8.3 software	71
Performing the initial configuration of storage subsystems on Solaris hosts	71
Adding storage subsystems to SMclient	72

Updating NVSRAM and FASSt firmware	72
Setting up a Solaris host group	73
Default partitioning for Solaris devices	74
Direct-attached and SAN-attached configurations	74
Creating a direct-attached configuration	74
Creating a SAN-attached configuration	75
Configuring JNI host bus adapter cards	76
Installing the JNI adapter driver package	76
Modifying the JNI HBA settings	77
Installing RDAC on Solaris hosts	82
Verifying external storage	85
Redistributing volumes in case of failure	85
Chapter 5. Completing the software installation.	87
Failover protection	87
Changing the host type	87
Understanding heterogeneous hosts	88
Configuring storage subsystems	88
Performing an initial automatic discovery of storage subsystems.	88
Configuring alert notifications.	90
Renaming storage subsystems	90
Creating a storage subsystem profile	90
Performing other storage subsystem management tasks	91
Chapter 6. Storage Manager 8.3 with high-availability cluster services.	93
General information	93
Using cluster services on HP-UX systems	93
Using cluster services on Solaris systems	93
General hardware requirements.	93
System dependencies	94
Using cluster services on AIX systems	95
General hardware requirements.	95
High Availability Cluster Management Program	95
Parallel System Support Programs and General Parallel File System	95
GPFS, PSSP, and HACMP cluster configuration diagrams	96
Appendix A. AIX system requirements	103
Hardware requirements	103
Software requirements	103
RDAC installation requirements	103
Appendix B. HP-UX system requirements	105
Appendix C. Solaris system requirements	107
Appendix D. MC/Service Guard configuration details	109
Appendix E. JNI host bus adapter settings	111
Notices	115
Trademarks.	115
Glossary	117
Index	125

Figures

1. Process flow by current publications	xii
2. Host-agent (in-band) managed storage subsystems	8
3. Direct (out-of-band) managed storage subsystems	10
4. Installation process (part 1)	11
5. Installation process (part 2)	12
6. Installation process (part 3)	13
7. Controller Properties window	34
8. Use of commands applied to validating system configuration.	35
9. Change Mappings window	42
10. Sample SMdevices command output (method 1)	59
11. Sample SMdevices command output (method 2)	60
12. Enterprise Management window	89
13. Cluster configuration with single FAStT storage server — one to four partitions	97
14. Cluster configuration with three FAStT storage servers — one partition per FAStT	98
15. Cluster configuration with four FAStT storage servers — one partition per FAStT	99
16. RVSD cluster configuration with two FAStT storage servers — two partitions per FAStT	100
17. HACMP/GPFS cluster configuration with one FAStT storage server — one partition.	101
18. HACMP/GPFS cluster configuration with two FAStT storage servers — two partitions per FAStT	102

Tables

1.	TotalStorage FAStT Storage Manager Version 8.3 titles by user tasks	xiii
2.	TotalStorage FAStT600 Fibre Channel Storage Server document titles by user tasks	xiv
3.	TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks	xv
4.	TotalStorage FAStT700 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.	Storage Partitioning premium feature	3
9.	Installation sequences of Storage Manager 8.3 software packages by host type	6
10.	FAStT storage servers that are supported on AIX systems	17
11.	Supported versions of hardware for AIX systems	18
12.	Attributes for dar devices	26
13.	Attributes for dac devices	27
14.	Attributes for hdisk devices	28
15.	FAStT Storage Server disk array device types	33
16.	FAStT storage servers that are supported on HP-UX systems	49
17.	Supported versions of hardware for HP-UX systems	50
18.	Sample record of logical drive preferred and alternate paths	61
19.	FAStT storage servers that are supported on Solaris systems	65
20.	Supported versions of hardware for Solaris systems	66
21.	JNI HBA model numbers, adapter drivers, and configuration files	76
22.	PSSP and GPFS software requirements	95
23.	Filesets required for AIX 4.3.3 RDAC	103
24.	Filesets required for AIX 5.1 RDAC	103
25.	Filesets required for AIX 5.2 RDAC (with APAR IY35693)	104
26.	HP-UX kernel parameter configuration requirements	106
27.	Configuration settings for FCI-1063	111
28.	Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562	111
29.	Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410	112
30.	Configuration settings for FC64-1063	112

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 redundant array of independent disks (RAID), small computer system interface (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. The arrows in the flow chart indicate the current publications that cover, in detail, each step in the installation process.

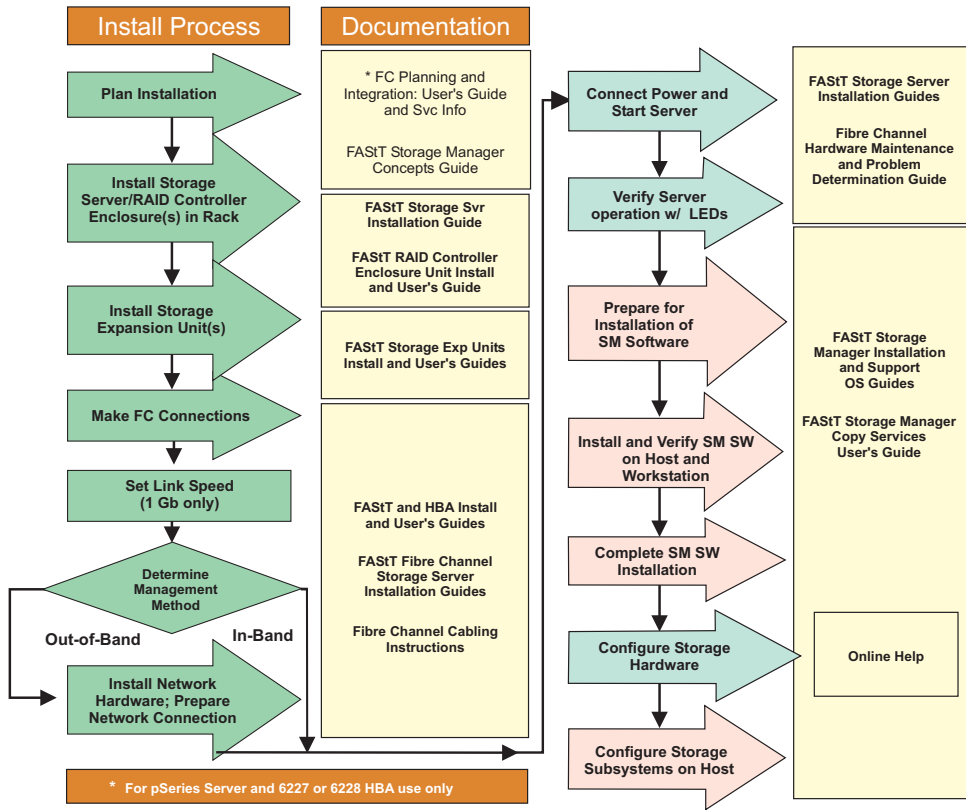


Figure 1. Process flow by current publications

FASTt documentation

The following tables present an overview of the FASTt Storage Manager and the FASTt600, FASTt900, FASTt700, 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.

FAStT Storage Manager Version 8.3 library

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

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

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
Installation and Support Guide for Windows® NT and Windows 2000, GC26-7522	✓		✓	✓		
Installation and Support Guide for Linux, GC26-7519	✓		✓	✓		
Installation and Support Guide for Novell NetWare, GC26-7520	✓		✓	✓		
Installation and Support Guide for UNIX and AIX Environments, GC26-7521	✓		✓	✓		
IBM TotalStorage FAStT Storage Manager Copy Services User's Guide, GC26-7561	✓		✓	✓	✓	
IBM FAStT Storage Manager Script Commands (see product CD)				✓		
IBM TotalStorage FAStT Storage Manager Concepts Guide, GC26-7560	✓	✓	✓	✓	✓	✓

FAStT600 Fibre Channel Storage Server library

Table 2 associates each document in the FAStT600 Fibre Channel Storage Server library with its related common user tasks.

Table 2. 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 FAStT600 Rack Mounting Instructions, 24P8125	✓	✓				
IBM TotalStorage FAStT600 Cabling Instructions, 24P8126	✓	✓				

FAStT900 Fibre Channel Storage Server library

Table 3 associates each document in the FAStT900 Fibre Channel Storage Server library with its related common user tasks.

Table 3. TotalStorage FAStT900 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
FAStT900 Installation and Support Guide, GC26-7530	✓	✓		✓		
FAStT900 Fibre Channel Cabling Instructions, 24P8135	✓	✓				
FAStT900 User's Guide, GC26-7534				✓	✓	✓
FAStT Host Adapter Installation and User's Guide, 59P5712		✓			✓	
FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
Fibre Channel Planning and Integration: User's Guide and Service Information, SC23-4329	✓	✓			✓	✓
FAStT Management Suite Java User's Guide, 32P0081					✓	✓
Fibre Channel Hardware Maintenance Manual and Problem Determination Guide, GC26-7528						✓

FAST700 Fibre Channel Storage Server library

Table 4 associates each document in the FAST700 Fibre Channel Storage Server library with its related common user tasks.

Table 4. TotalStorage FAST700 Fibre Channel Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
FAST700 Installation and Support Guide, 32P0171	✓	✓		✓		
FAST700 Fibre Channel Cabling Instructions, 32P0343	✓	✓				
FAST700 Fibre Channel Storage Server User's Guide, 32P0341				✓	✓	✓
EXP700 Storage Expansion Unit Installation and User's Guide, 32P0178	✓	✓		✓	✓	✓
FAST Host Adapter Installation and User's Guide, 59P5712		✓			✓	
FAST FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
TotalStorage FAST FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
FAST Management Suite Java User's Guide, 32P0081					✓	✓
Fibre Channel Hardware Maintenance Manual, 19K6130						✓
Fibre Channel Problem Determination Guide, 48P9804						✓

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
FAStT500 RAID Controller Enclosure Unit User's Guide, 48P9847				✓	✓	✓
FAStT EXP500 Storage Expansion Unit Installation and User's Guide, 59P5637	✓	✓		✓	✓	✓
FAStT Host Adapter Installation and User's Guide, 59P5712		✓			✓	
FAStT FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
TotalStorage FAStT FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
FAStT Management Suite Java User's Guide, 32P0081					✓	✓
Fibre Channel Hardware Maintenance Manual, 19K6130						✓
Fibre Channel Problem Determination Guide, 48P9804						✓

FAST200 Fibre Channel Storage Server library

Table 6 associates each document in the FAST200 Fibre Channel Storage Server library with its related common user tasks.

Table 6. TotalStorage FAST200 and FAST High Availability Storage Server document titles by user tasks

Title	User Tasks					
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance
FAST200 and FAST200 HA Storage Servers Installation and User's Guide, 59P6243	✓	✓		✓	✓	
FAST200 Fibre Channel Cabling Instructions, 21P9094	✓	✓				
FAST Host Adapter Installation and User's Guide, 59P5712		✓			✓	
FAST FC2-133 Dual Port Host Bus Adapter Installation and User's Guide, GC26-7532		✓			✓	
FAST FC2-133 Host Bus Adapter Installation and User's Guide, 48P9823		✓			✓	
FAST Management Suite Java User's Guide, 32P0081					✓	✓
Fibre Channel Hardware Maintenance Manual, 19K6130						✓
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 FASTt500 RAID Controller Enclosure Unit Installation Guide, 59P6244		✓			✓	
IBM FASTt500 RAID Controller Enclosure Unit User's Reference, 48P9847		✓			✓	
IBM Netfinity® Fibre Channel Cabling Instructions, 19K0906		✓				
IBM FASTt200 and FASTt200 HA Storage Servers Installation and User's Guide, 59P6243		✓			✓	
IBM FASTt200 Fibre Channel Cabling Instructions, 21P9094		✓				
IBM TotalStorage FASTt EXP700 Storage Expansion Unit Installation and User's Guide, 32P0178		✓			✓	
IBM FASTt EXP500 Installation and User's Guide, 59P5637		✓			✓	
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 17 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 49 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 65 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 87 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 93 provides information about high-availability clustering system options.

Appendix A, “AIX system requirements,” on page 103, Appendix B, “HP-UX system requirements,” on page 105, and Appendix C, “Solaris system requirements,” on page 107 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 109 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 111 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 17 and Appendix A, “AIX system requirements,” on page 103.
- For HP-UX, see Chapter 3, “Installing storage management station software on HP-UX systems,” on page 49 and Appendix B, “HP-UX system requirements,” on page 105.
- For Solaris, see Chapter 4, “Installing storage management station software on Solaris systems,” on page 65 and Appendix C, “Solaris system requirements,” on page 107.

Web sites

The most up-to-date information about FAS*St* Fibre Channel Storage Servers and FAS*St* Storage Manager 8.3, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following Web sites.

FAS*St* Fibre Channel Storage Servers

Find the latest information about all of the IBM FAS*St* Fibre Channel Storage Servers at:

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

IBM TotalStorage products

Find information about all IBM TotalStorage products at:

www.storage.ibm.com/

TotalStorage FAS*St* interoperability matrix

Find the latest information about operating system and host bus adapter support, clustering support, storage area network (SAN) fabric support, and Storage Manager feature support at:

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

Storage Manager 8.3 readme files

Find the latest readme files for FAS*St* Storage Manager 8.3 at:

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

Click the link for your FAS*St* Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAS*St* Storage Manager for your operating system.

Fix delivery center for AIX

Find the latest AIX downloads at:

techsupport.services.ibm.com/server/aix.fdc

FAS*St* technical support

Find FAS*St* downloads, hints and tips, documentation, parts information, HBA and fibre channel support at:

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

Premium feature activation

Enable a premium feature on a FAS*St*600 or FAS*St*900 by using the online tool at:

www.storage.ibm.com/pfeatures.html/

IBM publications center

Find IBM publications at:

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

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 (1-408-256-0488) 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. FTFA/Bldg. 050-1
5600 Cottle Road
San Jose, CA 95193-0001
U.S.A.

Chapter 1. Introduction

This document provides you 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 files

Read these first.

For the most recent readme files, see the following Web site:

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

Click the link for your FASTT Storage Server.

When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for your operating system.

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

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 Storage Manager 8.3 installation CD.

Overview of Storage Manager 8.3

The IBM TotalStorage FASTT Storage Manager Version 8.3 is a set of client and host tools that enable you to manage IBM FASTT200, FASTT500, FASTT600, FASTT700, and FASTT900 Fibre Channel Storage Server subsystems from a storage management station.

The Storage Manager 8.3 software is available on the product CD. You can also download Storage Manager 8.3 from the following Web site:

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

Click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for your operating system.

The storage management station

The *storage management station* is the system that is responsible for managing all, or a portion of, a storage network. It 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.

A storage management station can be either of the following configurations:

- A remote system, connected to an Ethernet network, that is used to manage one or more storage subsystems
- A host that is connected to the storage subsystem with a fibre-channel input/output (I/O) path, which manages the attached storage subsystems

Note: 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 host-agent communicates with the controller over the fibre-channel connection through the *access volume*.

Storage Manager 8.3 online help

After you have completed all the installation and configuration procedures that are provided in this document, refer to the following online help systems. The online help contains 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.

Changes in this edition

The following information has been changed from the previous edition of this document:

- New procedure for installing the RDAC driver on AIX
- New procedure for disabling cache mirroring on AIX
- Additional AIX dar and dac attributes, and corresponding AIX output examples
- Updated procedures for identifying logical drives on AIX
- Updated information about storage partitioning and setting up host groups
- Updated information about cluster services
- FlashCopy section removed; this information is now covered in the *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide*
- General editorial improvements

Storage Manager 8.3 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.

Premium copy services features

This document lists the premium copy services features that are supported by Storage Manager 8.3, but does not describe them in detail. For detailed copy services concepts and procedures, see the following documents:

- *IBM TotalStorage FAStT Storage Manager Copy Services User's Guide*
- *IBM TotalStorage FAStT Storage Manager Concepts Guide*

FlashCopy

The FlashCopy premium feature supports creating and managing

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.

Note: The Remote Mirror Option is not supported on the FAStT200 and FAStT600 storage servers. FlashCopy is not supported on the FAStT600 base model but is supported on the FAStT600 with the Turbo option. See “FAStT600 restrictions” on page 4 for more information.

Storage Partitioning premium feature

The Storage Partitioning feature enables you to associate a set of logical drives on a FAStT storage server that can only be accessed by specified hosts and host ports. This association of logical drives to a set of hosts and host ports is called a *storage partition*. The benefit of defining storage partitions is to allow controlled access to the logical drives to only those hosts that are also defined in the storage partition.

Table 8 shows whether storage partitioning is a standard feature on your FAStT storage server.

Table 8. Storage Partitioning premium feature

FAStT Storage Server	Storage Partitioning enabled by default	Maximum number of storage partitions
FAStT900	Yes (64 partitions standard)	64
FAStT700	Yes	64
FAStT600	No	16
FAStT500	Yes	64
FAStT200	Yes (16 partitions standard)	16

If you do not use Storage Partitioning, all logical drives are contained in a default host group. Logical drives in the default host group can be accessed by any fibre-channel initiator that has access to the FAStT host port.

When the FAStT storage server is attached to a SAN, you can use zoning within the fabric to limit access to the FAStT host ports to specific set of hosts. Also, when homogeneous host servers are directly attached to the FAStT storage server, access to all logical drives might be satisfactory.

Recommendation: On systems that have the Storage Partitioning feature enabled, always use storage partitioning when configuring logical drives and hosts. Configure each server to use separate partitions, unless your system uses supported clustering software applications.

Enabling your premium features

Complete the following procedures to enable a premium feature.

Before you begin: Gather the following information:

- Your Feature Activation Code, which accompanies the 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. Start SMclient. The Enterprise Management window opens.
2. In the left pane, double-click the storage subsystem for which you want to enable the premium feature. 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 the premium feature: Go to the following Web site:

www.storage.ibm.com/pfeatures.html/

Follow the instructions on the Web site to generate and download the feature key file, and enable the premium feature.

Your system is now ready to use the premium feature.

If the premium feature becomes disabled, you can access the Web site and repeat this process. For any additional assistance, please contact your IBM technical support representative.

FAStT600 restrictions

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

- FlashCopy
- Remote Mirror Option

The following restrictions also apply to the FAStT600:

- Supports 2 Gb host bus adapters (HBAs) only
- Supports 2 Gb link speed only
- Supports FAStT EXP700 expansion units only (maximum of three)
- Supports only one drive side arbitrated loop
- Supports only 42 drives
- Supports a maximum of 16 storage partitions

Note: If you want the ability to use the FlashCopy premium feature, additional storage partitions and drives, 1 Gb interface, and FAStT EXP100 expansion units, you can purchase the FAStT600 Turbo feature. The Turbo feature also includes FAStT Storage Manager Version 8.4.

Overview of Storage Manager 8.3 software packages

Storage Manager 8.3 contains the following client software packages:

SMruntime Storage Manager 8.3 Java™ compiler
SMclient Storage Manager 8.3 client package

Storage Manager 8.3 contains the following host software packages:

SMagent Storage Manager 8.3 agent package
SMutil Storage Manager 8.3 utility package
RDAC Storage Manager 8.3 multipath device drivers

Note: In addition to this document, see the *IBM FAStT Storage Manager Concepts Guide* for more information about Storage Manager 8.3 software.

SMruntime software package

The Storage Manager 8.3 runtime software, SMruntime, is a Java compiler for the Storage Manager 8.3 client software, SMclient. SMruntime must be installed before SMclient is installed.

SMclient software package

The Storage Manager 8.3 client software, SMclient, is a Java-based graphical user interface (GUI).

SMclient enables you to configure, manage, and troubleshoot the following devices in a storage subsystem, either through a host system or through a storage management station:

- IBM FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 Fibre Channel Storage Servers
- FAStT EXP500 and EXP700 Fibre Channel expansion units

Specifically, SMclient allows you to perform the following tasks:

- 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 (NVSRAM)
- View the major events log (MEL)
- Assign redundant RAID controllers

- Expand the storage capacity using FAStT EXP500 and EXP700 expansion units

Note: The FAStT600 supports EXP700 expansion units only. See “FAStT600 restrictions” on page 4 for more information.

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.

SMagent software package

The Storage Manager 8.3 agent (SMagent) package contains the optional host-agent software, which you can use on AIX, 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” on page 7.

SMutil 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 RDAC.

RDAC monitors I/O paths. If a component failure occurs in one of the fibre-channel paths, 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 24 to download them from the appropriate Web site.

Software installation sequence

Install the software packages in the sequences shown in Table 9.

Table 9. Installation sequences of Storage Manager 8.3 software packages by host type

Step	AIX	HP-UX	Solaris
1	SMruntime (see “Installing SMruntime” on page 20)	SMruntime (see “Installing the client software” on page 50)	SMruntime (see “Installing the client software” on page 67)

Table 9. Installation sequences of Storage Manager 8.3 software packages by host type (continued)

Step	AIX	HP-UX	Solaris
2	SMclient (see “Installing SMclient” on page 20)	SMclient (see “Installing the client software” on page 50)	SMclient (see “Installing SMclient” on page 68)
3	RDAC (see “Installing the RDAC driver” on page 24)	SMagent (see “Installing the SMagent (optional)” on page 52)	SMagent (see “Installing SMagent (optional)” on page 70)
4		SMutil (see “Installing SMutil (required)” on page 53)	SMutil (see “Installing SMutil (required)” on page 70)
5			RDAC (see “Installing RDAC on Solaris hosts” on page 82)

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:

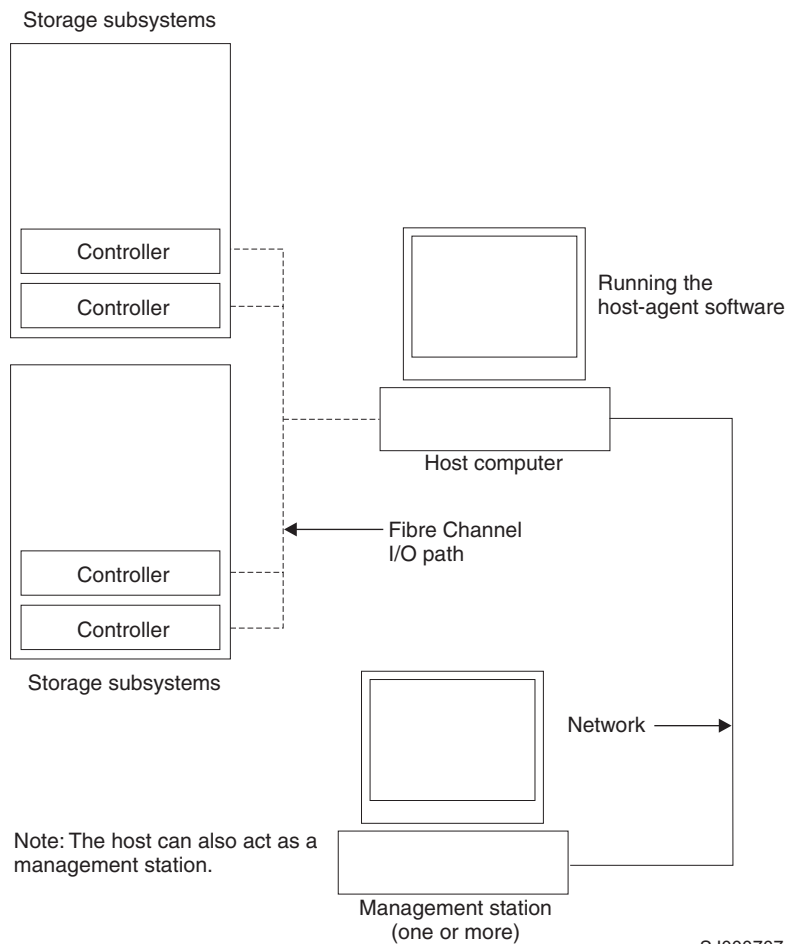
- If both fibre-channel connections to the controllers are lost, the SMclient cannot be accessed for problem determination.
- Host-agent (in-band) management is not supported on AIX.
- The direct (out-of-band) management method. In this method, you manage the storage subsystems directly over the network 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 8 shows the host-agent (in-band) management method.

Restrictions:

- If both fibre channel connections to the controllers are lost, the SMclient cannot be accessed for problem determination.
- Host-agent (in-band) management is not supported on AIX.



SJ000707

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

Host-agent (in-band) management method advantages

Managing storage subsystems by 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 17, Chapter 3, “Installing storage management station software on HP-UX systems,” on page 49, or Chapter 4, “Installing storage management station software on Solaris systems,” on page 65).
- 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.

Host-agent (in-band) management method disadvantages

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

- If both fibre channel connections to the controllers are lost, the SMclient cannot be accessed for problem determination.

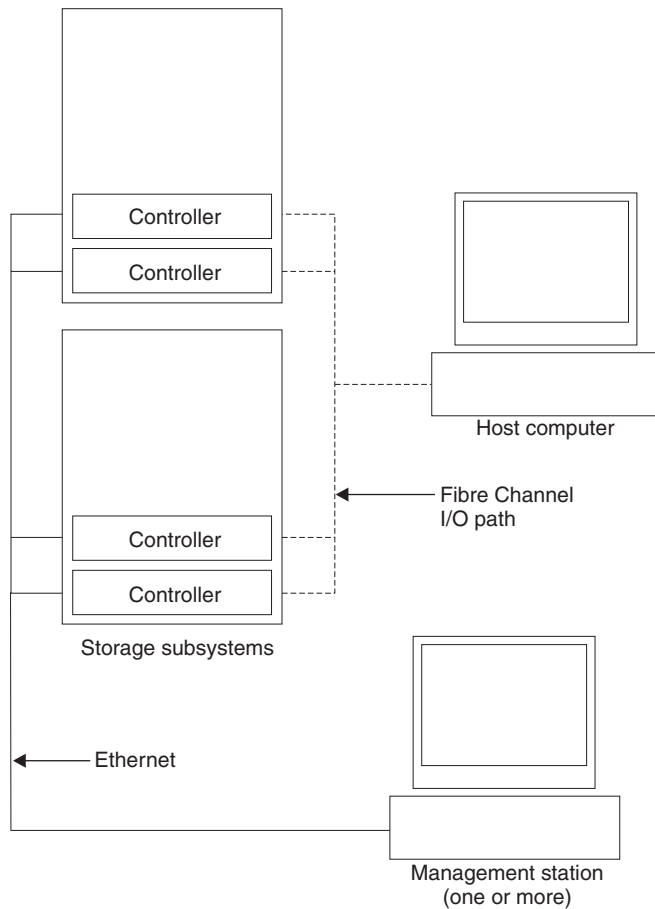
- Unlike direct (out-of-band) management, in which the fibre channel connections are dedicated for I/O, host-agent (in-band) management uses the fibre channel for both I/O and status reporting. Therefore, host-agent management uses more bandwidth on the fibre channel than direct management, and this can have an impact on performance.
- 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 on page 10 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.



SJ000708

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

Direct (out-of-band) management method advantages

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

- Unlike host-agent (in-band) management, which uses the fibre channel connections for both I/O and status reporting, direct (out-of-band) management dedicates the fibre channel for I/O only. Therefore, direct management can result in higher performance than host-agent management.
- Because status reporting is conducted over Ethernet connections, the SMClient can be accessed for problem determination if both fibre channel connections to the controllers are lost.
- The Ethernet connections to the controllers enable a storage management station running SMClient to manage storage subsystems that are connected to a host.

Direct (out-of-band) method disadvantages

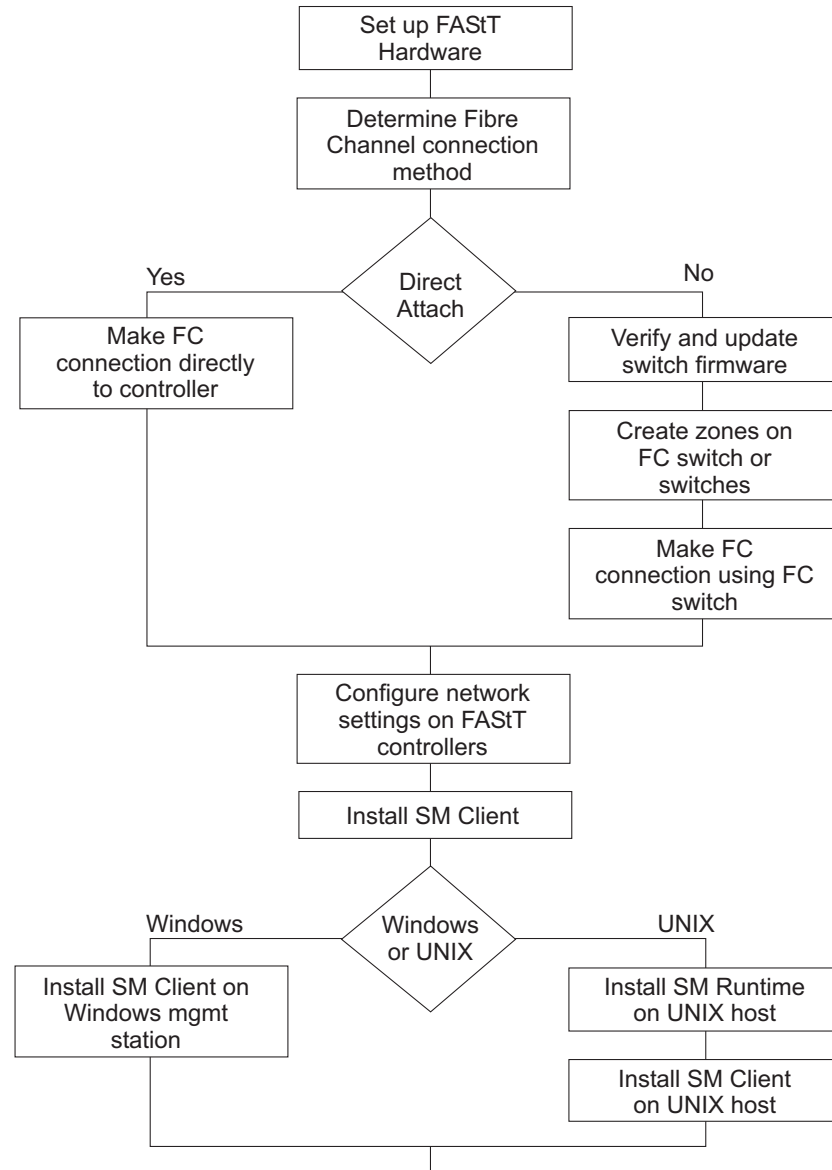
Managing storage subsystems by using the direct (out-of-band) management method has the following disadvantages:

- 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.

Installation process diagram

Figure 4, Figure 5 on page 12, and Figure 6 on page 13 provide an overview of the installation process for out-of-band management.

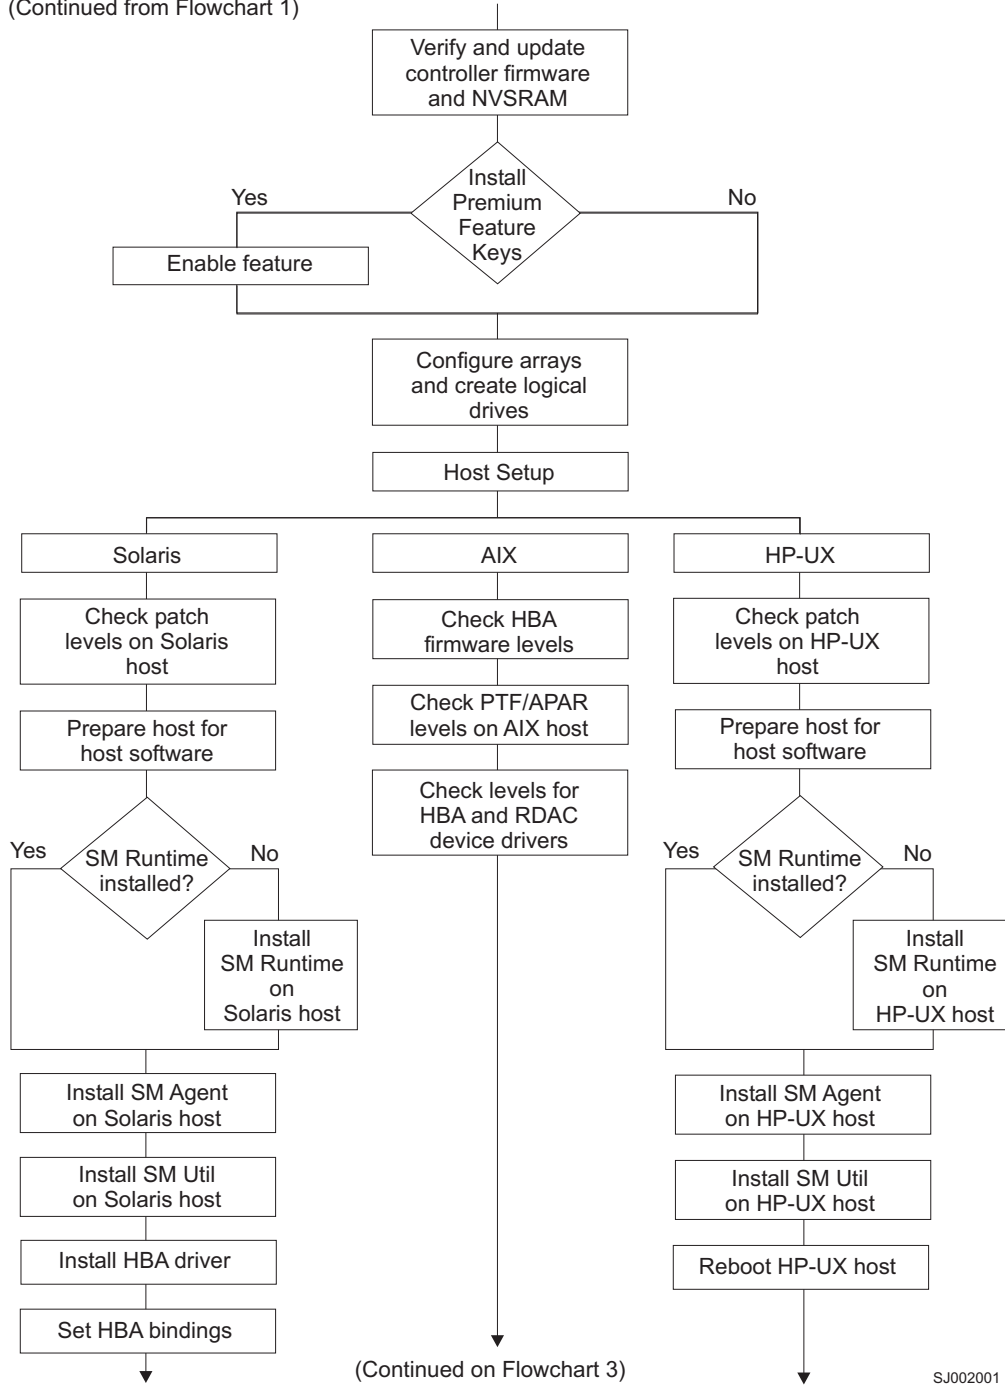


(Continued on Flowchart 2)

SJ002000

Figure 4. Installation process (part 1)

(Continued from Flowchart 1)

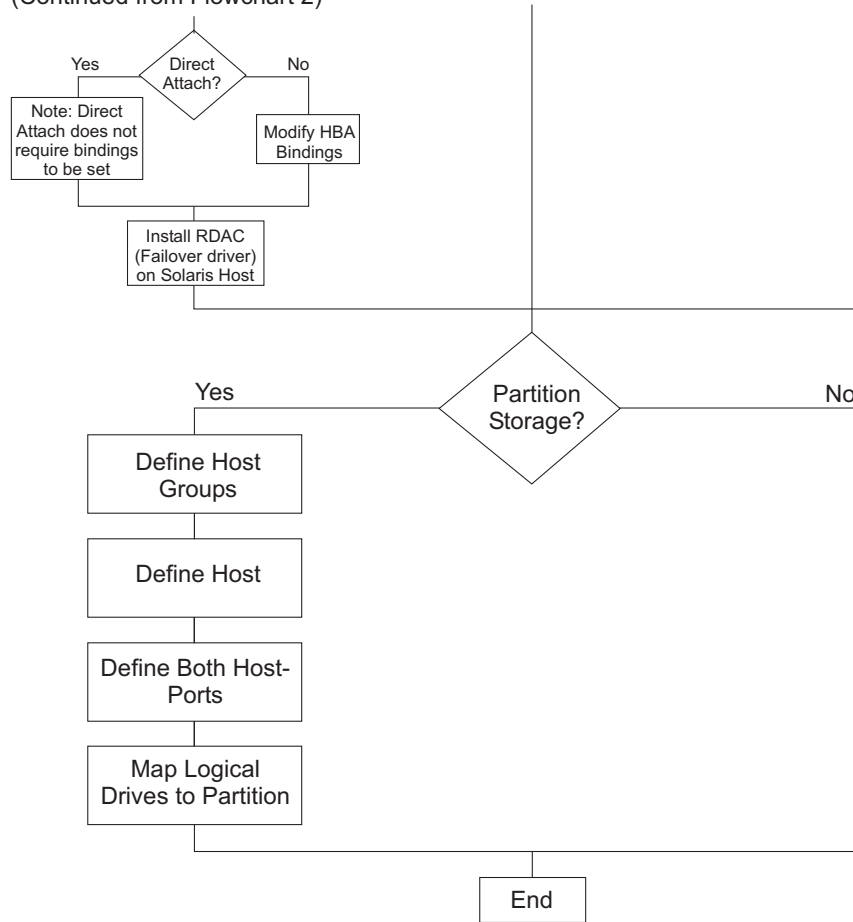


(Continued on Flowchart 3)

SJ002001

Figure 5. Installation process (part 2)

(Continued from Flowchart 2)



SJ002002

Figure 6. Installation process (part 3)

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 103
HP-UX versions: 11.0 (32 bit or 64 bit) 11.11 11i	FAStT200, FAStT500, FAStT600, FAStT700, FAStT900	Appendix B, "HP-UX system requirements," on page 105

Operating systems	FAStT FC Storage Servers	Where to find more information
Sun Solaris versions: 6 7 8	FAStT200, FAStT500, FAStT600, FAStT700, FAStT900	Appendix C, "Solaris system requirements," on page 107

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 by 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 step 1 through step 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.

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 storage area network (SAN)-attached configurations
- Installing RDAC
- AIX configuration information
- Implementing dynamic volume expansion (DVE) and dynamic capacity expansion (DCE)
- Common disk array errors

Hardware and firmware requirements

Table 10 lists the FAStT storage servers that are supported with Storage Manager 8.3.

Table 10. FAStT storage servers that are supported on AIX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 RAID Controller Enclosure Unit	1742-90U	Appware 5.30.12.00 NVSRAM N1742F900R830V05
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 5.30.12.00 NVSRAM N1742F700R830V04
IBM TotalStorage FAStT600 RAID Controller Enclosure Unit	1722-60U, 1722-60X	Snapware 5.33.07.00 NVSRAM N1722F600R833V03
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 5.30.12.00 NVSRAM N3552F500R830V04
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Snapware 5.30.12.00 NVSRAM N3542-2RUR830V06
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.12.00 NVSRAM N3542-2RUR830V06

Table 11 on page 18 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 11 on page 18. Note that checkmarks indicate which versions are supported on the FAStT600.

Table 11. 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 9324	✓
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9166	
IBM HA Emulex LP7000	FC 6227	S1F3.22A1 or higher	
IBM HA Emulex LP9000	FC 6228	C1D3.82A1 or higher	✓
IBM HA Emulex LP9802-F2	FC 6239	H1D1.00X5	✓
Brocade switch	2109-S08 2109-S16	2.6.1.j	
Brocade switch	2109-F16	3.1.1.b	✓
Brocade switch	2109-M12, 2109-F32	4.1.1.a	✓
Brocade switch	2109-F08	3.0.2.k	✓
McData switch	2032-064, 2032-140	4.01.00	✓
McData switch	2031-016	4.01.00	
McData switch	2031-216, 2031-224	4.01.00	✓
InRange switch	2042-001	4.3.0	✓
Cisco switch	2062-D01, 2062-D07, 2062-T07	1.1	✓

Note: Versions of firmware and NVSRAM that ship with this product might be later releases than those listed in this document. To ensure that you have the latest versions of the firmware and the NVSRAM, be sure to read the readme file that is shipped with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your FAStT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAStT Storage Manager for AIX.

AIX restrictions

The following restrictions apply to FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 storage servers:

- F-RAID Manager is not supported.
- Host-agent (in-band) management is not supported on AIX.
- The maximum number of partitions per AIX host, per FAStT storage server, is two.
- AIX hosts (or LPARs) can support multiple HBAs and FAStT devices. However, there is a restriction on the number of HBAs that can be connected to a single FAStT storage server. You can configure up to two HBAs per partition and up to

two partitions per FAStT storage server. Additional HBAs can be added to support additional FAStT storage servers and other SAN devices, up to the limits of your specific server platform.

- Direct-attach configurations are restricted to single-initiator configurations only. Only one connection to each FAStT minihub is allowed.

Note: FAStT200 and FAStT600 storage servers do not have minihubs.

- Single HBA configurations are allowed, but each single HBA configuration requires that both controllers in the FAStT be connected to the HBA, either through a switch or daisy-chain (direct-attach only). If they are connected through a switch, both controllers must be within the same SAN zone as the HBA.

Note: Having a single HBA configuration can lead to loss of access data in the event of a path failure.

- 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 20.

Exception: If you are running a FAStT600 configuration without partitioning enabled, you can use the default host group.

- Other storage devices, such as tape devices or other disk storage, must be connected through separate HBAs and SAN zones.
- Booting up 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 FAStT must be up and operational.
- Path failover is not supported during the AIX boot process. Once the AIX host has started, failover operates normally.
- Dynamic volume expansion (DVE) is supported only on AIX 5.2.

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 17 and “Software requirements” on page 103.
- 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 24 to download them from the appropriate Web site.

- For a list of AIX 4.3.3 filesets, see Table 23 on page 103.
- For a list of AIX 5.1 filesets, see Table 24 on page 103.
- For a list of AIX 5.2 filesets, see Table 25 on page 104.

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 /path_name  
/SMruntime.aix-08.30.G5.00.bff SMruntime.aix.rte
```

where *path_name* is the directory path to the SMruntime package.

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

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 /path_name/  
SMclient.aix-08.33.G5.03.bff SMclient.aix.rte
```

where *path_name* is the directory path to the SMclient package.

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 premium feature is enabled. For more information, see “Enabling your premium features” on page 4.

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 FASSt storage controllers” on page 14 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.

Upgrading FASSt firmware and NVSRAM

Before you begin: 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 and the NVSRAM, read the readme file that is shipped with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your FASSt Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASSt Storage Manager for AIX.

If you do not have the latest versions of the firmware and NVSRAM, click the **Back** button on your browser to return to the previous Web page. You can download the correct versions by clicking the appropriate link that is listed in the IBM FASSt Controller Firmware matrix.

Upgrading FASSt firmware

Complete the following steps to upgrade the controller firmware:

1. Ensure that all I/O to the controllers has been stopped.
2. Download the correct version of the firmware from the Web site. Place the file in a designated directory on the host system.
3. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **Firmware**.
4. Type or select the full path name of the firmware file.
5. Click **OK** to update both controllers.
6. Click **Yes** to start the download. A new window opens.
7. Close the Subsystem Management window and then reopen it to complete the firmware update.
8. Verify that the firmware was successfully installed.

Upgrading NVSRAM

Complete the following steps to upgrade the NVSRAM:

1. Ensure that all I/O to the controllers has been stopped.
2. Download the correct version of NVSRAM from the Web site. Place the file in a designated directory on the host system.
3. In the Subsystem Management window, click **Storage Subsystem** → **Download** → **NVSRAM**.
4. Type or select the full path name of the NVSRAM directory.
5. Double-click the NVSRAM filename or click **OK** to select the correct file.
6. Click **Yes** to start the download.
7. Verify that the NVSRAM was successfully installed.

Setting up an AIX host group

A *host group* is an entity in the Storage Partitioning topology that defines a logical collection of host computers that require shared access to one or more logical drives. Independently of the host group, you can grant individual hosts in a defined host group access to storage partitions. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 3.

Notes:

1. Create the AIX host group at the storage subsystem level. Do not create host groups at the default group level
Exception: If you are running a FASSt600 configuration without partitioning enabled, you can use the default host group.
2. 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.

Complete 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, complete 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 FAStT.
3. Type the following command:

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

where *fcs#* is the fcs number of the HBA that is connected to the FAStT. 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 minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.

Note: The FAStT200 and FAStT600 storage servers do not have minihubs.

- 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 minihub 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. See “FAStT600 restrictions” on page 4 for more information.

Installing the RDAC driver

Complete the following steps to update the RDAC driver version (devices.fcp.disk.array.rte) on an AIX system. Repeat these steps for all AIX systems that are connected to the storage subsystem.

Before you begin: Check the fileset versions that are listed in Appendix A, “AIX system requirements,” on page 103 or in the readme file, to verify whether your AIX systems have the most recent RDAC filesets. You only need to perform this installation on AIX systems that do not have the most recent RDAC filesets.

1. Download the most recent filesets from the following Web site:

techsupport.services.ibm.com/server/aix.fdc

2. 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 IBM technical support representative. If it does not return an error, then you are finished installing the updated RDAC driver on this AIX system.

3. 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 FAST storage subsystem. The RDAC driver creates the following devices that represent the FAST storage subsystem configuration:

- dar** The disk array router represents the entire array, including the current and the deferred paths to all LUNs (hdisk 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.
- utm** The universal transport mechanism (utm) device is used only in host-agent (in-band) configurations, as a communication channel between the SMagent and the FAST.

Note: You might see the utm device listed in command output even though in-band configuration is not supported on AIX. For example, a utm might be listed when you run the **lsattr** command on a dac.

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 31.

See “Viewing Object Data Manager attributes in AIX” on page 30 for information about using the **lsattr** command to view attribute settings on an AIX system.

Important: You might discover that the configuration process has created two dacs and two dars on one FAStT subsystem. This situation can occur when your host is using a partition that does not have any associated LUNs. When that happens, the system cannot associate the two dacs under the correct dar. If there are no LUNs, the system generates two dacs as expected, but it also generates two dars.

The following list shows the most common causes:

- You create a partition and attach the LUNs to it, but you do not add the host ports to the partition. Therefore, the host ports remain in the default partition.
- You replace one or more HBAs, but do not update the worldwide name (WWN) of the partition for the HBA.
- You switch the FAStT from one set of HBAs to another as part of a reconfiguration, and do not update the WWNs.

After you resolve the problem and run **cfgmgr** again, in each of these cases, the system removes the extra dar or moves it from the Available state to the Defined state. After the system moves the dar into the Defined state, you can delete it.

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 12, Table 13 on page 27, and Table 14 on page 28. Display the default with the **lsattr -dl** command.

Table 12. 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.

Table 12. Attributes for dar devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<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>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.	True	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 FAST 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 FAST.
<i>switch_retries</i>	Number that specifies how many times to retry failed switches, in integers.	True	0 - 255. Should not be changed.

Table 13. 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 FAST configuration.

Table 13. Attributes for dac devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<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.
<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>node_name</i>	Name of the fibre-channel node.	False	Set by FAStT.
<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 14. 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.

Table 14. Attributes for hdisk devices (continued)

Attribute	Definition	Changeable (T/F)	Possible value
<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
<i>reserve_lock</i>	Indicator that shows whether the fcparray 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 14. 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 attributes in AIX

To view the Object Data Manager (ODM) 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** parameter 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
aen_freq 600 Polled AEN frequency in seconds True
all_controller dac0,dac1 Available Controllers False
autorecovery no Autorecover after failure is corrected True
balance_freq 600 Dynamic Load Balancing frequency in seconds True
cache_size 128 Cache size for both controllers False
fast_write_ok yes Fast Write available False
held_in_reset none Held-in-reset controller True
hlthchk_freq 600 Health check frequency in seconds True
load_balancing no Dynamic Load Balancing True
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
GLM_type low GLM type False
alt_held_reset no Alternate held in reset False
cache_size 128 Cache Size in MBytes False
controller_SN 1T24594458 Controller serial number False
ctrl_type 1722-600 Controller Type False
location Location Label True
lun_id 0x0 Logical Unit Number False
node_name 0x200200a0b80f14af FC Node Name False
passive_control no Passive controller False
scsi_id 0x11000 SCSI ID False
utm_lun_id 0x001f000000000000 Logical Unit Number False
ww_name 0x200200a0b80f14b0 World Wide Name False
```

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

```
lsattr -El hdisk174
cache_method fast_write Write Caching method False
ieee_volname 600A0B80000F4365000001063F7076A7 IEEE Unique volume name False
lun_id 0x0069000000000000 Logical Unit Number False
prefetch_mult 12 Multiple of blocks to prefetch on read False
pvid none Physical volume identifier False
q_type simple Queuing Type False
queue_depth 2 Queue Depth True
raid_level 5 RAID Level False
reassign_to 120 Reassign Timeout value True
reserve_lock yes RESERVE device on open True
rw_timeout 30 Read/Write Timeout value True
scsi_id 0x11f00 SCSI ID False
size 2048 Size in Mbytes False
write_cache yes Write Caching enabled False
```

Changing ODM attribute settings in AIX

As discussed previously, ODM attributes for the RDAC driver and FAS*T* 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.

Disabling FASTT cache mirroring: Cache mirroring is enabled on the FASTT storage server by default. However, you might want to disable this function in certain environments where high performance is a critical factor (such as streaming media applications).

For most applications, cache mirroring should remain enabled.

Attention: You must be careful when you disable cache mirroring, particularly when the write cache function is enabled.

- If a controller failure or a power down occurs, data that is in the write cache (but not written to the physical disk media) might be lost. This can result in corrupted files, file systems, or databases.
- On AIX 5.1 and 5.2, cache mirroring is *not* automatically re-enabled the next time you reboot the system or run **cfgmgr**.

In the SMclient, complete the following steps to disable cache mirroring.

Before you begin: If write cache is enabled, make backups of all data before disabling cache mirroring.

1. In the Subsystem Management window, right-click the volume on which you want to disable cache mirroring, and select **Change —> Cache Settings**.
2. In the Change Cache Settings window, clear the *Enable write caching with mirroring* checkbox.
3. Click **OK**.

Notes:

1. When a LUN is opened that is running with write cache enabled and cache mirroring disabled, an FCP array warning message displays. The warning displays again every 24 hours until cache mirroring is enabled again. See FCP_ARRAY_ERR25, which is described in “Resolving disk array errors” on page 45.
2. If you are running AIX 4.3.3, do not disable cache mirroring while cache write is enabled. If this condition exists, the RDAC software automatically re-enables it the next time you reboot the system or run **cfgmgr**.

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 file systems and system panics.

Verifying the installation and configuration of AIX hosts

After you have completed the procedures that are described in the previous sections, 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 the FASTT volume, as shown in Table 15:

Table 15. FASTT Storage Server disk array device types

FASTT Storage Server	Disk array device type
FASTT200	3542 (200)
FASTT500	3552 (500)
FASTT600	1722 (600)
FASTT700	1742 (700)
FASTT900	1742-900

The following example shows the results of the command for a set of FASTT900 LUNs:

```
# lsdev -Cc disk
hdisk0 Available 10-88-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk32 Available 31-08-01 1742-900 Disk Array Device
hdisk33 Available 91-08-01 1742-900 Disk Array Device
hdisk34 Available 31-08-01 1742-900 Disk Array Device
hdisk35 Available 91-08-01 1742-900 Disk Array Device
```

Identifying the controller ID numbers

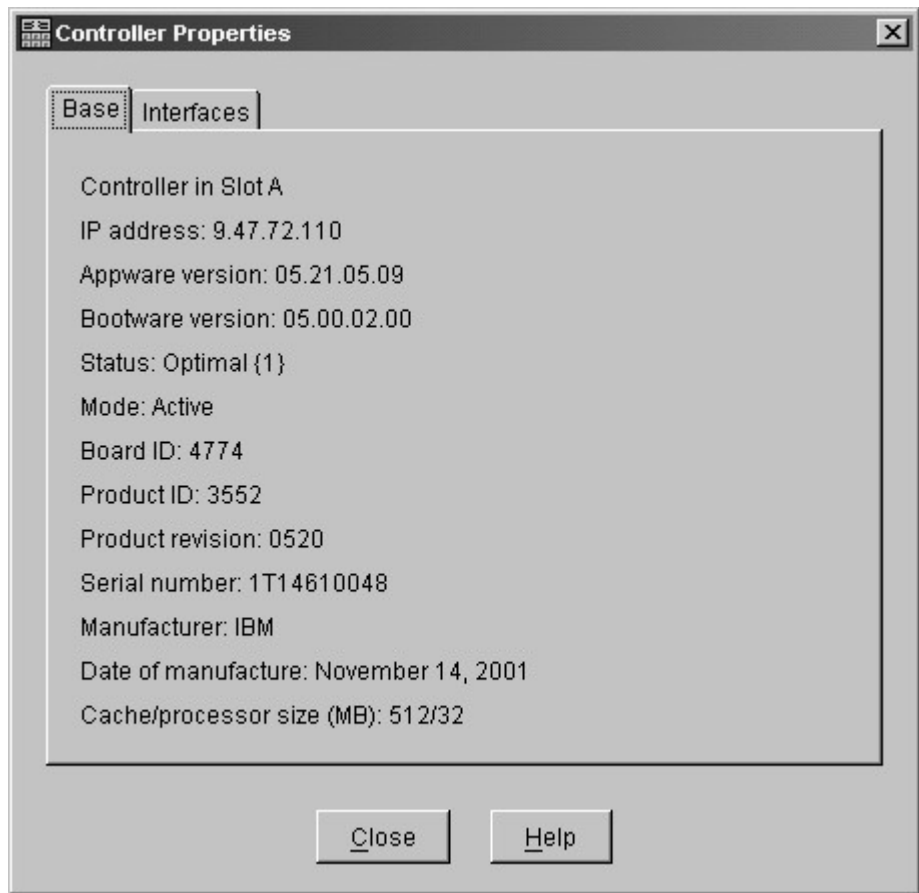
One of the attributes listed by the **lsattr -EI** command is the controller serial number (`controller_SN`) of that dac.

Complete the following step to display the SMclient Controller Properties window:

In the Logical/Physical view of the SMclient, right-click a controller and select **Properties**. The Controller Properties window opens.

Figure 7 shows the Controller Properties window of the controller in slot A. This controller has a serial number of 1T14610048 and is represented by dac1.

```
# lsattr -El dac1
GLM_type          low                GLM type          False
alt_held_reset    no                 Alternate held in reset False
cache_size        1024             Cache Size in MBytes False
controller_SN    1T14610048      Controller serial number False
ctrl_type         1742-0900        Controller Type    False
location          location          Location Label     True
lun_id            0x0              Logical Unit Number False
node_name         0x200400a0b80f434d FC Node Name      False
passive_control   no                Passive controller False
scsi_id           0x11a00          SCSI ID           False
utm_lun_id        none              Logical Unit Number False
ww_name           0x200500a0b80f434e World Wide Name   False
```



SJ000709

Figure 7. 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 previously 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 FAST 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 FAST (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 FAST storage subsystems. The use of commands in validating configurations is illustrated in Figure 8.

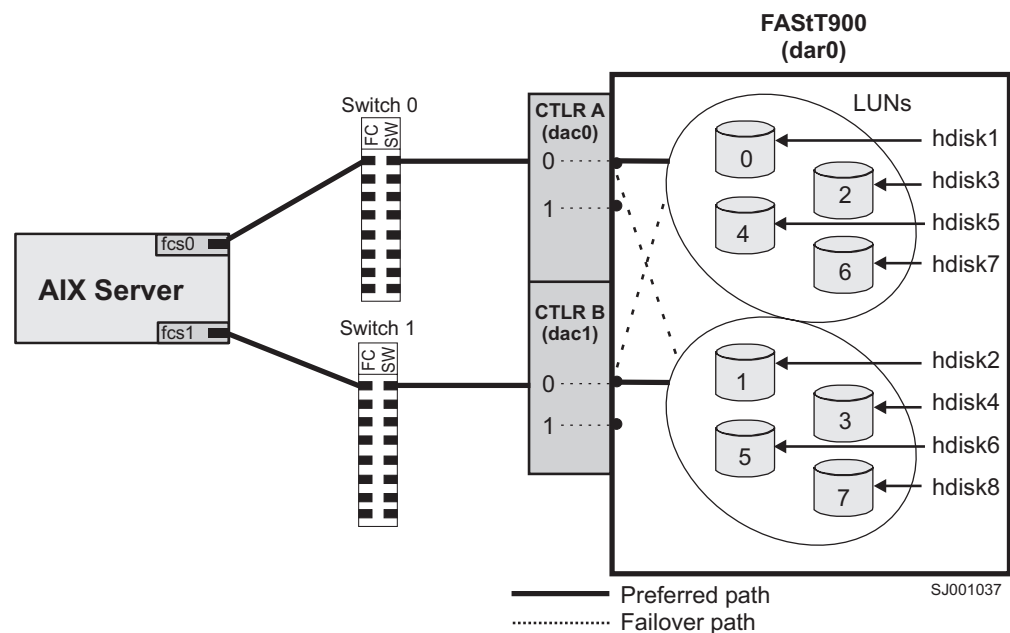


Figure 8. Use of commands applied to validating system configuration

Using the lsdev command

The following example uses the **lsdev** command to show the status of the dar, which represents a FAST storage subsystem. This example shows dar as a machine type 3542, which is a FAST200. 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 FAST 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. Note 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 shown 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

You can use the **fget_config** command to display current configuration information for a FAStT subsystem.

Use the following command syntax for the **fget_config** command:

```
fget_config -A | -l | -v dar_name
```

- A Displays output for all the configured dars in the subsystem. If you use this parameter, do not specify a dar name.

- **l** Displays output only for the dar that you specify.
- **v** Displays more information about the dar or dars, such as the user array names that were created when you configured the FASTT subsystem.

The following example shows a sample output from the **fget_config -l dar3** command. A dac can be Active, Passive, or Offline. Only dacs that are in Active state are used for I/O. In the output, you can see that dar3 consists of dac6 and dac9.

```
# fget_config -l dar3

dac6 ACTIVE dac9 ACTIVE
hdisk26 dac9
hdisk27 dac6
hdisk28 dac9
hdisk29 dac6
hdisk30 dac9
hdisk31 dac6
hdisk32 dac9
hdisk33 dac9
hdisk83 dac6
```

The following example shows a sample output from the **fget_config -vl dar3** command. Note that using the **-v** parameter causes the user array name and the LUN names to be displayed:

```
# fget_config -vl dar3

User array name = 'D59 FASTt700 #1'
dac6 ACTIVE dac9 ACTIVE

Disk    DAC    LUN Logical Drive
utm          10
hdisk26 dac9    1 multipart11
hdisk27 dac6    2 multipart12
hdisk28 dac9    3 multipart13
hdisk29 dac6    4 multipart14
hdisk30 dac9    5 multipart15
hdisk31 dac6    6 multipart16
hdisk32 dac9    7 multipart17
hdisk33 dac9    9 multipart19
hdisk83 dac6    8 multipart18
```

The following example shows a sample output from the **fget_config -A** command. Note that using the **-A** parameter causes all the dars to be displayed:

```

# fget_config -A

---dar0---
dac3 ACTIVE dac0 PASSIVE
hdisk4 dac3
hdisk5 dac3
hdisk6 dac3
hdisk7 dac3
hdisk8 dac3
hdisk10 dac3
hdisk11 dac3
hdisk12 dac3
hdisk13 dac3
hdisk14 dac3

---dar1---
dac1 ACTIVE dac4 ACTIVE
hdisk15 dac1
hdisk16 dac1
hdisk17 dac4
hdisk18 dac1
hdisk19 dac4
hdisk20 dac1
hdisk21 dac4
hdisk22 dac1
hdisk23 dac4
hdisk24 dac1

---dar2---
dac2 ACTIVE dac5 ACTIVE
hdisk25 dac5

---dar3---
dac6 ACTIVE dac9 ACTIVE
hdisk26 dac9
hdisk27 dac6
hdisk28 dac9
hdisk29 dac6
hdisk30 dac9
hdisk31 dac6
hdisk32 dac9
hdisk33 dac9
hdisk83 dac6

```

The following example shows a sample output from the **fget_config -Av** command. Note that using the **-v** parameter causes the user array names and the LUN names to display:

```

# fget_config -Av

---dar0---
User array name = 'D59S Flute-2
(Fast-T500) #1'
dac3 ACTIVE dac0 PASSIVE

Disk   DAC   LUN Logical Drive
hdisk4 dac3   0 dualpart0
hdisk5 dac3   1 dualpart1
hdisk6 dac3   2 dualpart2
hdisk7 dac3   4 dualpart4
hdisk8 dac3   5 dualpart5
hdisk10 dac3  8 dualpart8
hdisk11 dac3  9 dualpart9
hdisk12 dac3 10 dualpart10
hdisk13 dac3 11 dualpart11
hdisk14 dac3 12 dualpart12

---dar1---
User array name = 'D59 FASTt700 #1'
dac1 ACTIVE dac4 ACTIVE

Disk   DAC   LUN Logical Drive
utm    10
hdisk15 dac1  0 multipart20
hdisk16 dac1  1 multipart1
hdisk17 dac4  2 multipart2
hdisk18 dac1  3 multipart3
hdisk19 dac4  4 multipart4
hdisk20 dac1  5 multipart5
hdisk21 dac4  6 multipart6
hdisk22 dac1  7 multipart7
hdisk23 dac4  8 multipart8
hdisk24 dac1  9 multipart9

---dar2---
User array name = 'D59S FASTt500 #2'
dac2 ACTIVE dac5 ACTIVE

Disk   DAC   LUN Logical Drive
utm    31
hdisk25 dac5  0 1

---dar3---
User array name = 'D59 FASTt700 #1'
dac6 ACTIVE dac9 ACTIVE

Disk   DAC   LUN Logical Drive
utm    10
hdisk26 dac9  1 multipart11
hdisk27 dac6  2 multipart12
hdisk28 dac9  3 multipart13
hdisk29 dac6  4 multipart14
hdisk30 dac9  5 multipart15
hdisk31 dac6  6 multipart16
hdisk32 dac9  7 multipart17
hdisk33 dac9  9 multipart19
hdisk83 dac6  8 multipart18

```

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
GLM_type          low          GLM type          False
alt_held_reset    no           Alternate held in reset False
cache_size        1024         Cache Size in MBytes False
controller_SN     1T23559428   Controller serial number False
ctrl_type         1742-0900    Controller Type    False
location          location      Location Label     True
lun_id            0x0          Logical Unit Number False
node_name         0x200400a0b80f434d FC Node Name      False
passive_control   no           Passive controller False
scsi_id          0x11100     SCSI ID           False
utm_lun_id        none         Logical Unit Number False
ww_name          0x200400a0b80f434e World Wide Name   False

# lsattr -El dac1
GLM_type          low          GLM type          False
alt_held_reset    no           Alternate held in reset False
cache_size        1024         Cache Size in MBytes False
controller_SN     1T23559493   Controller serial number False
ctrl_type         1742-0900    Controller Type    False
location          location      Location Label     True
lun_id            0x0          Logical Unit Number False
node_name         0x200400a0b80f434d FC Node Name      False
passive_control   no           Passive controller False
scsi_id          0x11a00     SCSI ID           False
utm_lun_id        none         Logical Unit Number False
ww_name          0x200500a0b80f434e World Wide Name   False

# lsattr -El hdisk1
----- Parts removed:
scsi_id          0x11100     SCSI ID           False
size              4096         Size in Mbytes    False
write_cache       yes          Write Caching enabled False

```

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 logical drive, including information that allows you to map the system device name to the logical drive on the FAStT storage subsystem.

The following example shows the output of an **lsattr** command for a LUN named hdisk4:


```

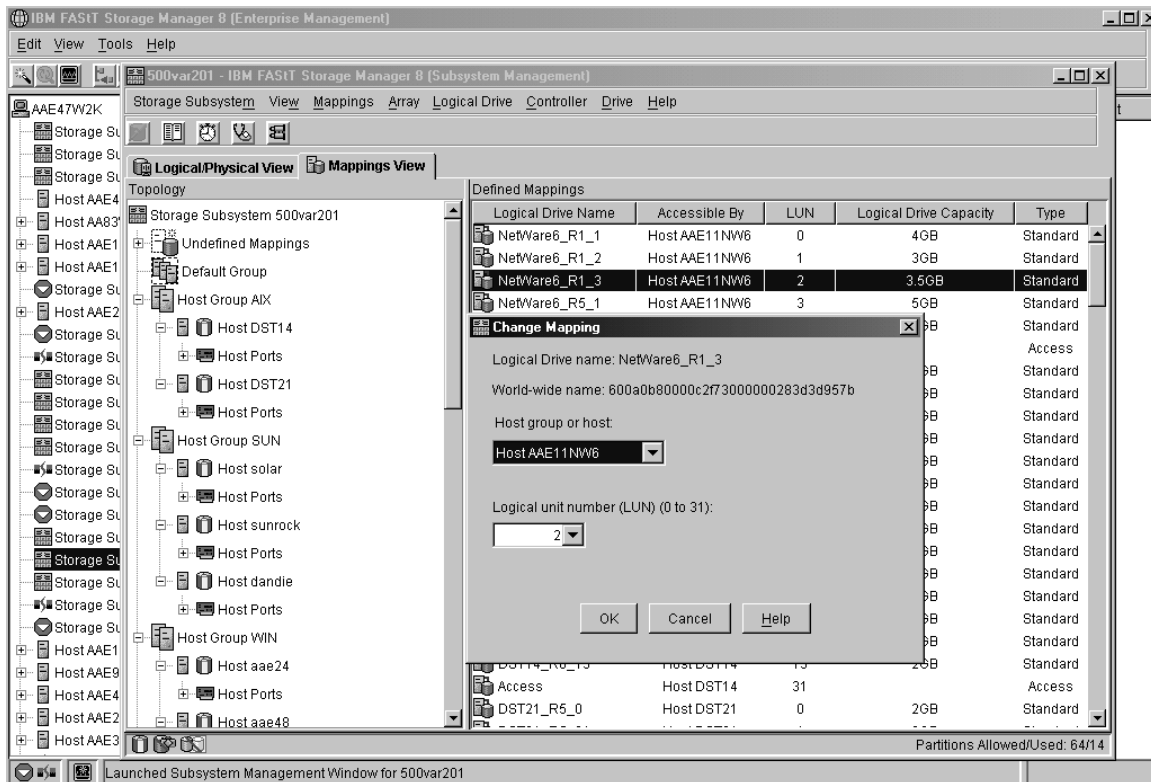
# lsattr -El hdisk4
pvid          none                Physical volume identifier      False
q_type        simple              Queuing Type                   False
queue_depth   30                          Queue Depth                    True
reserve_lock  yes                          RESERVE device on open        True
write_cache   yes                          Write Caching enabled         True
size         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      0x0002000000000000        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 600A0B80000C2F73000000283D3D957B IEEE Unique volume name      False

```

In the example, you can see that hdisk4 is a 6 GB LUN of type RAID 0, with a LUN ID of 2, and an IEEE volume name of 600A0B80000C2F73000000283D3D957B.

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 value of the worldwide name on the FASTt storage subsystem. You can see the worldwide name of this LUN in the Change Mapping window of the SMclient, which is shown in Figure 9 on page 42.

To view the Change Mapping window, click the **Mappings view** tab in the Subsystem Management window. Right-click the name of the logical drive whose attributes you want to see, and select **Change Mapping**.



SJ000712

Figure 9. Change Mappings window

You can also identify the system disk name by viewing the *Subsystem ID (SSID)* attribute on the Logical Drive Properties window of the SMclient. The value of the *Subsystem ID (SSID)* attribute on the FASTt storage subsystem is the same as the value of the *lun_id* attribute on the AIX host. To open the Logical Drive Properties window, right-click the name of the logical drive in the Subsystem Management window and select **Properties**.

Identifying FlashCopy volume types

Complete 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 displays, as shown in the following example:

```
[root@e1m17c171] / # 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 displays, as shown in the following example:

```
[root@elm17c171] / # fget_config -Av
---dar0---
User array name = 'FAST 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 34.

Using fast I/O failure for fibre-channel devices

I/O failover occurs when a fibre-channel adapter driver detects a link failure on the fabric. You can change the failover characteristics by setting the fcscli device attribute **fc_err_recov** to one of the following settings.

Note: Fast I/O failure is supported only on AIX 5.1 and 5.2.

fast_fail Enables fast I/O failure.

If the fibre-channel adapter driver detects a link loss, it waits a few seconds to allow the fabric to stabilize. If **fast_fail** is specified, the adapter then begins failing all I/O at the adapter driver if it detects that the device is not on the fabric. Any new I/O or future retries of the failed I/O are failed immediately.

Fast I/O failure can be useful in multipath configurations. It can decrease the I/O fail times due to link loss between the storage device and the switch, and can allow faster failover to alternate paths. It does not, however, change the failover timing between the host HBA and the switch.

delayed_fail Default setting.

If **delayed_fail** is specified, I/O failure proceeds as normal; retries are not immediately failed, and failover takes longer than it does if **fast_fail** is specified.

In single-path configurations, especially configurations with a single path to a paging device, use the **delayed_fail** setting.

Example: You can enable fast I/O failure by setting this attribute, as shown in the following example. Be sure to stop all I/O and put the fcscli device into a Defined state before you set the attribute.

```
chdev -l fcscli0 -a fc_err_recov=fast_fail
```

Notes:

1. The **fast_fail** attribute only affects failover that occurs between the switch and the FAStT storage server. It does not affect failover that occurs between the host and the switch.
2. Set the **fast_fail** attribute on each HBA that is configured to the FAStT storage server.
3. You can use fast I/O failure only in a SAN environment. You cannot use it in a direct-attach environment.

Using dynamic tracking of fibre-channel devices

AIX supports dynamic tracking of fibre-channel devices, which allows the dynamic movement of a fibre-channel path between the fabric switch and the FAStT subsystem by suspending I/O for 15 seconds while the move occurs.

Note: Dynamic tracking is supported only on AIX 5.2.

You can enable or disable dynamic tracking by setting the fscsi device attribute **dyntrk** to one of the following settings:

yes Enables dynamic tracking.

If dynamic tracking is enabled, the fibre-channel adapter detects when the fibre-channel node port ID of a device changes. It reroutes the traffic that is destined for that device to the new WWPN while the devices are still online.

For example, you can move a cable from one switch port to another while the devices are still online, and no failover occurs if you complete the move within 15 seconds. After 15 seconds, failover occurs.

Note: The ports must be in the same zone on the same switch.

no Default setting.

If dynamic tracking is not enabled, you must take the devices offline before you move a cable from one port to another. Otherwise, failover occurs.

Example: You can enable dynamic tracking by setting this attribute, as shown in the following example. Be sure to stop all I/O and put the fscsi device into a Defined state before you set the attribute.

```
chdev -l fscsi0 -a dyntrk=yes
```

Notes:

1. Set the dynamic tracking attribute on both HBAs on the controller.
2. You can use dynamic tracking only in a SAN environment. You cannot use it in a direct-attach environment.

Using dynamic capacity expansion and dynamic volume expansion

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

Note: DVE is only supported on AIX 5.2.

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, complete 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 DCE:

1. From the Logical/Physical view of the SMclient, right-click **Array** → **Add Free Capacity (Drives)**.
2. From the Logical/Physical view of the SMclient, right-click **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 FAStT is 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 either reboot the system or run the **cfgmgr** command before you can view the additional capacity using the ODM attribute.

Performing a dynamic volume expansion operation

Complete the following steps to perform a DVE on the AIX host. Refer to the man page for the **chvg** command for more information about expanding LVM volumes.

1. Unmount the file systems from the LVM volume group that has been expanded by typing the following command:

```
chvg -g 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. You can import the volume group to AIX 5.1 (or lower).

2. Mount the file systems.

Notes:

1. You might have to run the **varyoffvg** command, then the **varyonvg** command, on the volume group for LVM to see the size change on the disk. However, the new capacity is available for use by the operating system.
2. You cannot re-size while the volume group is activated in classic or enhanced concurrent mode.
3. Re-sizing for the root volume group is not supported.

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.

Note: For more information about troubleshooting, see the *IBM TotalStorage FAStT Fibre Channel Hardware Maintenance Manual and Problem Determination Guide*.

- **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*
An undetermined error has occurred.
- **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*
One array controller has become unavailable, so I/O has moved to the other controller.
- **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 might 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*
The controllers are unable to communicate with each other. This could result from one of the controllers being rebooted while the error log was being generated. However, it could be a much more serious error that indicates a problem with the fibre-channel connections.
- **FCP_ARRAY_ERR14** *ARRAY DRIVE FAILURE*
A serious or unrecoverable error has been detected on a physical disk within the FAStT subsystem. A system engineer might be able to obtain the exact cause from an analysis of the sense data.
- **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 error log. 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.

- **FCP_ARRAY_ERR19** *SNAPSHOT VOLUME'S REPOSITORY FULL*

The repository capacity limit has been reached. To resolve this error you can increase the repository capacity.

- **FCP_ARRAY_ERR20** *SNAPSHOT OPERATION STOPPED BY ADMIN*

The FlashCopy (snapshot) operation has been disabled or stopped. To resolve this error you can re-create the FlashCopy.

- **FCP_ARRAY_ERR21** *SNAPSHOT REPOSITORY METADATA ERROR*

There was a problem with the metadata of the FlashCopy (snapshot) repository during the FlashCopy operation. To resolve this error you can re-create the FlashCopy.

- **FCP_ARRAY_ERR22** *REMOTE VOL MIRRORING: ILLEGAL I/O ORIGIN*

The primary volume received I/O from a remote array, or the secondary volume received I/O from other than the primary volume. To resolve this error you can try the operation again.

- **FCP_ARRAY_ERR23** *SNAPSHOT OPERATION NOT ALLOWED*

The repository capacity limit has been reached, so the FlashCopy (snapshot) operation has failed. To resolve this error you can delete or re-create the FlashCopy.

- **FCP_ARRAY_ERR24** *SNAPSHOT VOLUME'S REPOSITORY FULL*

The repository capacity limit has been reached. To resolve this error you can delete or re-create the FlashCopy (snapshot).

- **FCP_ARRAY_ERR25** *CACHED DATA WILL BE LOST IF CONTROLLER FAILS*

This message is a warning that a disk array volume (LUN) is running with write cache enabled and cache mirroring disabled. The warning displays when the LUN is opened, and it displays again every 24 hours until cache mirroring is enabled again.

If a controller failure or a power down occurs while the LUN is running in this mode, data that is in the write cache (but not written to the physical disk media) might be lost. This can result in corrupted files, file systems, or databases.

Redistributing volumes in case of failure

If you enabled autorecovery on the AIX host, you do not need to redistribute volumes manually after a controller failover. However, if you have a heterogeneous environment, you might need to redistribute logical volumes manually. Hosts that do not support some form of autorecovery, or AIX hosts that have autorecovery disabled, will not automatically redirect volumes to the preferred paths.

Complete 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
- Redistributing volumes in case of failure

Hardware and firmware requirements

For system-level hardware requirements, software requirements, and patch levels, see Appendix B, “HP-UX system requirements,” on page 105.

Attention: Be sure to install the following patch for your HP-UX operating system:

- **HP-UX 11.0:** PHKL_27003, version 1.0
- **HP-UX 11.i:** PHKL_26519, version 1.0

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

Table 16. FAStT storage servers that are supported on HP-UX systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 RAID Controller Enclosure Unit	1742-90U	Appware 5.30.12.00 NVSRAM N1742F900R830V05
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 5.30.12.00 NVSRAM N1742F700R830V04
IBM TotalStorage FAStT600 RAID Controller Enclosure Unit	1722-60U, 1722-60X	Snapware 5.33.07.00 NVSRAM N1722F600R833V03
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 5.30.12.00 NVSRAM N3552F500R830V04
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Snapware 5.30.12.00 NVSRAM N3542-2RUR830V06
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.12.00 NVSRAM N3542-2RUR830V06

Table 17 on page 50 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 17 on page 50. Note that checkmarks indicate which versions are supported on the FAStT600.

Table 17. Supported versions of hardware for HP-UX systems

Product Name	Model	Product release and firmware version	Supported on FAStT600
IBM TotalStorage FAStT700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9324	✓
IBM FAStT500 EXP500 Storage Expansion Unit	3560-1RU	ESM 9166	
HP FC Adapter	A5158A	B.11.00.06	
HP FC Adapter 2Gb	A6795A	B.11.00.01	✓
Brocade switch	2109-S08 2109-S16	2.6.0j	
Brocade switch	2109-F16	3.0.2n	✓
Brocade switch	2109-M12, 2109-F32	4.0.2b	✓
Brocade switch	2109-F08	3.0.2n	✓
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	1.30	✓

Note: Versions of firmware and NVSRAM that ship with this product might be later releases than those listed in this document. To ensure that you have the latest versions of the firmware and the NVSRAM, be sure to read the readme file that is shipped with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your FAStT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAStT Storage Manager for HP-UX.

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 105.
- Neither SMruntime nor SMclient is installed on the host and you are ready to install them on this machine.

Installing SMruntime

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 /path_name /cdrom
# swinstall -s /cdrom/HP-UX/SMruntime_0804500
```

where *path_name* is the directory path to the SMruntime package.
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

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 (required)

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 105.
- 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.

Adjust these instructions as required for you 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 (required)

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. 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 88 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. To ensure that you have the latest versions of the firmware and the NVSRAM, read the readme file that is shipped with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your FAStT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAStT Storage Manager for HP-UX.

Note: If you do not have the latest versions of firmware and NVSRAM, download them and complete the procedure described in “Updating NVSRAM and FAStT firmware” on page 55. If the version numbers are current, you can skip that procedure.

- If you have purchased a Storage Partitioning Premium Feature, ensure that the feature is enabled. See “Enabling your premium features” on page 4 for more information.

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

Adding storage subsystems to SMclient

Complete the following steps to add 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.

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

A *host group* is an entity in the Storage Partitioning topology that defines a logical collection of host computers that require shared access to one or more logical drives. Independently of the host group, you can grant individual hosts in a defined host group access to storage partitions. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 3.

Notes:

1. Create the HP-UX host group at the storage subsystem level. Do not create host groups at the default group level
Exception: If you are running a FAStT600 configuration without partitioning enabled, you can use the default host group.
2. 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.

Before you begin: 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
```

Complete the following steps to set up an HP-UX 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, 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 FAStT200, FAStT500, FAStT600, FAStT700, and FAStT900 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 FAStT storage server:
 - FAStT200 can support one HP-UX server.
 - FAStT500, FAStT600, FAStT700, and FAStT900 can support two AIX servers.
- Two-server FAStT500, FAStT700 or FAStT900 configurations require four host-side minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.

Note: The FAStT200 and FAStT600 storage servers do not have minihubs.

- There are two or four HBAs (HP A6795A or HP A5158A) 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, complete the following steps:

1. Connect the HBAs to each controller or minihub 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 a 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, complete the following 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.

Important:

- There are two methods for establishing redundant access to storage using PV-links:

- If you have firmware 5.xx.xx.xx installed, use method 1.
- If you have firmware 4.xx.xx.xx installed, use method 2.
- For both methods, you must have SMutil installed on the host.

Using PV-links: Method 1

If you have firmware 5.xx.xx.xx installed, use the following procedure to enable multipath I/O 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 10 on page 59. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each RAID controller has one logical drive access. For example, a subsystem that contains two RAID controllers has two logical drive accesses.

```
#SMdevices
```

```

/dev/rds/c166t0d0 [Storage Subsystem FAStT, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]
/dev/rds/c166t0d1 [Storage Subsystem FAStT, Logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Alternate Path (Controller-B): Not In Use]
/dev/rds/c166t0d2 [Storage Subsystem FAStT, Logical Drive Finance,
LUN 2, Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Alternate Path (Controller-B): Not In Use]
/dev/rds/c166t0d3 [Storage Subsystem FAStT, Logical Drive Purchasing,
LUN 3, Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Alternate Path (Controller-B): Not In Use]
/dev/rds/c166t0d4 [Storage Subsystem FAStT, Logical Drive Development,
LUN 4, Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>,
Preferred Path (Controller-B): In Use]
/dev/rds/c166t3d7 [Storage Subsystem FAStT, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f56d00000001b00000000>]

/dev/rds/c172t0d0 [Storage Subsystem FAStT, Logical Drive Accounting, LUN 0,
Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Alternate Path (Controller-A): Not In Use]
/dev/rds/c172t0d1 [Storage Subsystem FAStT, logical Drive HR, LUN 1,
Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>,
Preferred Path (Controller-A): In Use]
/dev/rds/c172t0d2 [Storage Subsystem FAStT, Logical Drive Finance, LUN 2,
Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>,
Preferred Path (Controller-A): In Use]
/dev/rds/c172t0d3 [Storage Subsystem FAStT, Logical Drive Purchasing, LUN 3,
Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>,
Preferred Path (Controller-A): In Use]
/dev/rds/c172t0d4 [Storage Subsystem FAStT, Logical Drive Development, LUN 4,
Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>,
Alternate Path (Controller-A): Not In Use]
/dev/rds/c172t3d7 [Storage Subsystem FAStT, Logical Drive Access, LUN 31,
Logical Drive WWN <600a0b80000f5d6c0000002200000000>]

```

Figure 10. Sample SMdevices command output (method 1)

Note: If you do not see the logical drives and logical drive accesses after running the **hot_add** and **SMdevices** commands, restart the HP-UX host by running the **reboot** command.

```
#reboot
```

- Determine the preferred and alternate path for each logical drive by examining the output from the **SMdevices** command, as shown in the example in Figure 10. Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path.

Preferred path

In the sample output that is shown in Figure 10, the preferred path is **/dev/rds/c166t0d0**:

```

/dev/rds/c166t0d0 [Storage Subsystem FAStT, Logical Drive
Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>,
Preferred Path (Controller-B): In Use]

```

Alternate path

In the sample output that is shown in Figure 10, the alternate path is **/dev/rds/c172t0d0**:

```
/dev/rdisk/c172t0d0 [Storage Subsystem FAStT, Logical Drive Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>, Alternate Path (Controller-A): NotIn Use]
```

Using PV-links: Method 2

If you have firmware 4.xx.xx.xx installed, use the following procedures to enable multipath I/O by using PV-links:

- Determine the preferred and alternate paths
- Create the volumes and volume groups

Determining preferred and alternate paths

Complete the following steps to determine the preferred and alternate paths:

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 11. Notice that every logical drive and logical drive access unit has been assigned a logical unit number (LUN). Each logical drive has two paths. Each RAID controller has one logical drive access. For example, a subsystem that contains two RAID controllers has two logical drive accesses.

```
#SMdevices
```

```
/dev/rdsk/c166t0d0 [Storage Subsystem FAStT, Logical Drive Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdsk/c166t0d1 [Storage Subsystem FAStT, Logical Drive HR, LUN 1, Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdsk/c166t0d2 [Storage Subsystem FAStT, Logical Drive Finance, LUN 2, Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdsk/c166t0d3 [Storage Subsystem FAStT, Logical Drive Purchasing, LUN 3, Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdsk/c166t0d4 [Storage Subsystem FAStT, Logical Drive Development, LUN 4, Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdsk/c166t3d7 [Storage Subsystem FAStT, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f56d00000001b0000000>]

/dev/rdsk/c172t0d0 [Storage Subsystem FAStT, Logical Drive Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
/dev/rdsk/c172t0d1 [Storage Subsystem FAStT, logical Drive HR, LUN 1, Logical Drive WWN <600a0b80000f5d6c000000273eaeae30>]
/dev/rdsk/c172t0d2 [Storage Subsystem FAStT, Logical Drive Finance, LUN 2, Logical Drive WWN <600a0b80000f5d6c000000253eaeadf8>]
/dev/rdsk/c172t0d3 [Storage Subsystem FAStT, Logical Drive Purchasing, LUN 3, Logical Drive WWN <600a0b80000f5d6c000000243eaeadbe>]
/dev/rdsk/c172t0d4 [Storage Subsystem FAStT, Logical Drive Development, LUN 4, Logical Drive WWN <600a0b80000f56d00000001d3eaeacef>]
/dev/rdsk/c172t3d7 [Storage Subsystem FAStT, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f5d6c000000220000000>]
```

Figure 11. Sample SMdevices command output (method 2)

Note: If you do not see the logical drives and logical drive accesses after running the **hot_add** and **SMdevices** commands, restart the HP-UX host by running the **reboot** command.

```
#reboot
```

- Determine the preferred and alternate path for each logical drive by examining the output from the **SMdevices** command, as shown in the example in Figure 11 on page 60.

Notice that each device is listed twice; one instance is the preferred path and one instance is the alternate path. Also, notice that each device has a worldwide name (WWN). 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 logical drive access in Figure 11 on page 60, you notice that they differ in only five digits, f56d0 and f5d6c.

The devices in Figure 11 on page 60 are viewed through the controllers c166 and c172. To determine the preferred path of a specific logical drive seen by the operating system, complete the following steps:

- Find the WWN for each logical drive access. In this case, Logical Drive Access 1 is associated with c166 and has the WWN of f56d0:

```
/dev/rdisk/c166t3d7 [Storage Subsystem FASTT, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f56d00000001b00000000>]
```

Logical Drive Access 2 is associated with c172 and has the WWN of f5d6c:

```
/dev/rdisk/c172t3d7 [Storage Subsystem FASTT, Logical Drive Access, LUN 31, Logical Drive WWN <600a0b80000f5d6c0000002200000000>]
```

- Identify the preferred device path name for the attached storage device by matching the logical drive WWN to a logical drive access WWN. In this case, the WWN for **LUN 0** is associated with controller c166 and c172. Therefore, the preferred path for **LUN 0** is /dev/rdisk/c166t0d0, which is controller c166:

```
/dev/rdisk/c166t0d0 [Storage Subsystem FASTT, Logical Drive Accounting, LUN 0, Logical Drive g<600a0b80000f56d00000001e3eaead2b>]
```

The alternate path is /dev/rdisk/c172t0d0, which is controller c172:

```
/dev/rdisk/c172t0d0 [Storage Subsystem FASTT, Logical Drive Accounting, LUN 0, Logical Drive WWN <600a0b80000f56d00000001e3eaead2b>]
```

- To keep a record for future reference, enter this path information for LUN 0 into a matrix (similar to the one in Table 18).

Table 18. Sample record of logical drive preferred and alternate paths

LUN	Logical Drive Name	Preferred Path	Alternate Path
0	Accounting	/dev/rdisk/c166t0d0	/dev/rdisk/c172t0d0
1	HR	/dev/rdisk/c172t0d1	/dev/rdisk/c166t0d1
2	Finance	dev/rdisk/c172t0d2	/dev/rdisk/c166t0d2
3	Purchasing	/dev/rdisk/c172t0d3	/dev/rdisk/c166t0d3
4	Development	/dev/rdisk/c166t0d4	/dev/rdisk/c172t0d4

- d. Repeat steps 3a on page 61 through 3c on page 61 for each logical drive that is seen by the operating system.

Continue to the next procedure.

Creating volumes and volume groups

After you have determined the preferred and alternate paths, and have recorded them in a matrix for future reference, complete the following steps to create volumes and volume groups.

Note: The steps in this procedure refer to LUN 0 in Table 18 on page 61.

1. Create a physical volume and define the primary paths for the attached storage devices. The primary path will be the preferred path. Type the following command at the shell prompt:

```
#pvcreate /dev/rdisk/c166t0d0
```

The system confirms the creation of the new physical volume.

2. Create volume groups.

Note: For more information on how to create volume groups, refer to HP-UX documentation or to man pages.

- a. Make a directory for volume group by typing the following commands. This directory must reside in the /dev directory.

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

- b. Create the group special file in the /dev directory for the volume group by typing the following command:

```
#mknod /dev/vg1/group c 64 0x010000
```

- c. Create a volume group and define physical volume names (primary link) for the attached storage device by typing the following command:

```
#vgcreate /dev/vg1/ /dev/dsk/c166t0d0
```

- d. Define the secondary path name (alternate path) for the attached-storage device by typing the following command:

```
#vgextend vg1 /dev/dsk/c172t0d0
```

Note: You can also use the **vgextend** command to add additional storage devices to an existing volume group. Add the primary path first, then add the alternate path, as shown in the following example:

- 1) Add the primary path for LUN1.

```
#vgextend vg1 /dev/dsk/c172t0d1
```

- 2) Add the secondary path for LUN1.

```
#vgextend vg1 /dev/dsk/c166t0d1
```

3. Create logical volumes. For more information, refer to HP-UX documentation.
4. Create file systems for the logical volumes.
5. Repeat steps 1 on page 62 through 4 to create additional volume groups. For more information, refer to HP-UX documentation.
6. Verify the primary (preferred) and secondary (alternate) paths for each device by typing the following command, where *vgname* is the volume group name:

```
#vdisplay -v vgname
```

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, complete the following steps to redistribute logical drives to their paths:

1. Repair or replace any faulty components. For more information, see the *IBM TotalStorage FAS*T* Installation and Support Guide* for the appropriate FAS*T* Fibre Channel Storage Server.
2. 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 information and procedures for Solaris operating systems:

- Hardware and firmware requirements
- Installing client software
- Installing host software
- Uninstalling Storage Manager 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 19 lists the FAStT storage servers that are supported with Storage Manager 8.3.

Table 19. FAStT storage servers that are supported on Solaris systems

Product Name	Model	Minimum appware/snapware (firmware) and NVSRAM versions
IBM TotalStorage FAStT900 RAID Controller Enclosure Unit	1742-90U	Appware 5.30.12.00 NVSRAM N1742F900R830V05
IBM TotalStorage FAStT700 RAID Controller Enclosure Unit	1742-1RU	Appware 5.30.12.00 NVSRAM N1742F700R830V04
IBM TotalStorage FAStT600 RAID Controller Enclosure Unit	1722-60U, 1722-60X	Snapware 5.33.07.00 NVSRAM N1722F600R833V03
IBM FAStT500 RAID Controller Enclosure Unit	3552-1RU	Appware 5.30.12.00 NVSRAM N3552F500R830V04
IBM FAStT200 RAID Controller Enclosure Unit	3542-1RU and FC 2101	Snapware 5.30.12.00 NVSRAM N3542-2RUR830V06
IBM FAStT200 RAID and Storage Unit, Double Controller	3542-2RU	Snapware 05.30.12.00 NVSRAM N3542-2RUR830V06

Table 20 on page 66 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 20 on page 66. Note that checkmarks indicate which versions are supported on the FAStT600.

Table 20. Supported versions of hardware for Solaris systems

Product Name	Model	Product release and firmware version	Supported on FASt600
IBM TotalStorage FASt700 EXP700 Storage Expansion Unit	1740-1RU	ESM 9324	✓
IBM FASt500 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	FCE2-1473 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	FCI-1063	fca-pci.2.5.18.pkg	
JNI FC Adapter	FCC-6562	5.2.1	
JNI FC Adapter	FCC2-6562	5.2.1	
JNI FC Adapter	FCX-6562	5.2.1	
JNI FC Adapter	FCC\X2-6562	5.2.1	
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 listed in this document. To ensure that you have the

latest versions of the firmware and the NVSRAM, be sure to read the readme file that is shipped with the product. You can also find the latest readme file at the following Web site:

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

Click the link for your FAS*T* Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAS*T* Storage Manager for Solaris.

For system-level hardware and software requirements and patch levels, see Appendix C, “Solaris system requirements,” on page 107.

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 6, 7, or 8 and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements,” on page 107.
- Neither SMruntime nor SMclient is installed on the host and you are ready to install them on this machine.

Installing SMruntime

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 65 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-file_name.pkg
```

where *file_name* is the name of the package.

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 6, 7, and 8.

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-file_name.pkg
```

where *file_name* is the name of the package.
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 68 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 (required)
- 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 82.

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 6, 7, or 8, and meets the minimum hardware and software requirements described in Appendix C, “Solaris system requirements,” on page 107.
- 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.

Complete 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-file_name.pkg
```

where *file_name* is the name of the package.

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 (required)

Complete the following steps to install SMutil:

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

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

where *path_name* is the path to the package and *file_name* is the name of the package.

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:

- Adding storage subsystems to SMclient
- Updating FAStT firmware and NVSRAM
- 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 your premium features” on page 4 for more information.

- 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 and the NVSRAM, read the readme file that is shipped with the product. You can also find the latest readme file at the Web site that is appropriate for your FASTT Storage Server:

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

Click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for Solaris.

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

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**.

Updating NVSRAM and FASTT firmware

Before you begin: 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 and the NVSRAM, read the readme file that is shipped with the product. You can also find the latest readme file at the Web site that is appropriate for your FASTT Storage Server:

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

Click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for Solaris.

If you do not have the latest versions of the firmware and NVSRAM, press the **Back** button on your browser to return to the previous Web page. You can download the correct versions by clicking the appropriate link that is listed in the IBM FASTT Controller Firmware matrix.

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.

Setting up a Solaris host group

A *host group* is an entity in the Storage Partitioning topology that defines a logical collection of host computers that require shared access to one or more logical drives. Independently of the host group, you can grant individual hosts in a defined host group access to storage partitions. You can make logical drive-to-LUN mappings to the host group or to an individual host in a host group.

For more information about storage partitioning, see “Storage Partitioning premium feature” on page 3.

Notes:

1. Create the Solaris host group at the storage subsystem level. Do not create host groups at the default group level.
Exception: If you are running a FAST600 configuration without partitioning enabled, you can use the default host group.
2. 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.

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

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:



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.

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 Solaris servers.

Requirements:

- Two-server FAStT500, FAStT700, or FAStT900 configurations require four host-side minihubs, each with exactly one fibre-channel connection from each HBA to a minihub.

Note: The FAStT200 and FAStT600 storage servers do not have minihubs.

- 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.

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

1. Connect the HBAs to each controller or minihub 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.

Complete 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.

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. See “FAStT600 restrictions” on page 4 for more information.

Configuring JNI host bus adapter cards

Before you can install the failover driver, you must install the HBA driver package and modify the HBA settings using 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.

Installing the JNI adapter driver package

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 FAStT 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.

The procedures for installing the JNI adapter driver package and modifying the JNI HBA settings both refer to Table 21.

Table 21. JNI HBA model numbers, adapter drivers, and configuration files

JNI HBA model numbers	JNI adapter driver package	Configuration file
FC64-1063-N	fca2.5.18.pkg	/kernel/drv/fca.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, FCE2-1473, FCE-6412 (2 Gb), FCE-6460	JNIC146x.pkg, version 5.2.1	/kernel/drv/jnic146x.conf

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

1. In Table 21, 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 drive package is installed:

```
#pkginfo adapter_driver_package
```

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

Modifying the 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 failover 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 82. You do not need to modify JNI HBA settings if you have a direct-attached configuration.

Modifying the JNI HBA settings for McData, Cisco, InRange, and Brocade 1 GB switches

If you have a Brocade 2 Gb switch, you need to force the HBA (FCC-6460, FCE-1473, or FCE-21473) to be a public loop device. See “Modifying the JNI HBA settings for a Brocade 2 Gb switch” on page 79 for the correct modification procedures for a Brocade 2 Gb switch.

Before you begin: Table 21 on page 76 lists the configuration file for each JNI adapter driver package. Also, some of the steps in these procedures refer to Appendix E, “JNI host bus adapter settings,” on page 111.

Complete the following procedures to modify the JNI HBA settings.

Modifying the loop settings:

1. 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
```

2. In the Vi Editor, uncomment and modify the loop attributes using the information in Appendix E, “JNI host bus adapter settings,” on page 111.
3. Run the following command to save changes made to the JNI configuration file.

```
#:wq
```

4. Run the following command to reboot the Solaris host.

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

Gathering the binding information:

1. After the host reboots, run the following commands to gather binding information in the messages log.

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

The messages log displays.

2. Search the messages log for the most recent HBA information. You can do this with a backwards search for the installed HBAs.

```
# /Sun0S
```

3. 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

Binding the HBAs to storage controllers: The examples in this procedure assume that you are binding two JNIC FCE-6460 HBAs to the FAStT controllers.

1. Run the following commands to open the JNI configuration file.

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

where *JNI_configuration_file* is the name of the JNI configuration file of the HBA that you want to modify.

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

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

2. Set and change the target variables to the specified values, as shown in the following examples.

target0_hba variables:

- a. Set the target to the first value that you recorded 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.

- b. Set the target to the second value that you recorded 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 variables:

- a. Set the target to the WWPN of the first controller using the value that you recorded from the */var/adm/messages* file.

Example:

```
target_Z_wwpn="controller_1_wwpn"
```

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

- b. Set the target to the WWPN of the second controller using the value that you recorded from the /var/adm/messages file.

Example:

```
target_Y_wwpn="controller_2_wwpn"
```

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

3. Run the following command to save the changes to the JNI configuration file.

```
#:wq
```

4. Run the following command to reboot the Solaris host.

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

Modifying the JNI HBA settings for a Brocade 2 Gb switch

If you are using a Brocade 2 Gb switch, you need to force the HBA (FCC-6460, FCE-1473, or FCE2-1473 operating in 2 Gb mode) to be a public loop device.

If you are using any other switch, see “Modifying the JNI HBA settings for McData, Cisco, InRange, and Brocade 1 GB switches” on page 77 for the correct modification procedures.

Complete the following procedures to modify the JNI HBA settings.

Modifying the loop settings:

1. Open the /var/adm/messages file by typing the following command:

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

2. Perform a backwards search of the /var/adm/messages file by typing the following command:

```
# /Sun0S
```

3. Look for the HBA in the /var/adm/messages file. The following example is for an FCE-6460-N adapter.

```

May 6 10:02:06 solar unix: jnic146x0:
Hba: JNI,FCR Model: FCE-6460-N
May 6 10:02:06 solar unix: jnic146x0:
FCode: Version 3.8.9 [ba79]
May 6 10:02:06 solar unix: jnic146x0:
IRQ: 1 Channel: 1
May 6 10:02:06 solar unix: jnic146x0:
SubVend: 1242 Sub: 6460 Rev: 100Vend:242 Dev: 1460
May 6 10:02:06 solar unix: jnic146x0:
WWNN: 100000017300284e WWPN: 200000017300284e
May 6 10:02:07 solar unix: jnic146x0:
Auto-negotiating link speed
May 6 10:02:07 solar unix: jnic146x0:
Configured as Public Loop port
May 6 10:02:08 solar unix: jnic146x0:
JNIC v5.2.1 (02111400)
May 6 10:02:08 solar unix: jnic146x0:
Copyright(c) 1995-2002 JNI Corp, All Rights Reserved.
May 6 10:02:10 solar unix: jnic146x0:
Link Up: 2Gig
May 6 10:02:11 solar unix: jnic146x0:
Port 011000 (WWN 200200a0b80f478e:200300a0b80f478f) online.
May 6 10:02:12 solar unix: jnic146x0: Network Port
fffc01 (WWN 1000006069500f0d:20fd006069500f0d) available

```

4. Record the target. In the following example, the target for the first FCE-6460-N adapter is target 0. The second FCE-6460-N adapter is target 3:

```

May 6 10:02:11 solar unix: jnic146x0: Port 011000
(WWN 200200a0b80f478e:200300a0b80f478f) online.

May 6 10:02:11 solar unix: jnic146x3: Port 011000
(WWN 200200a0b80f478e:200300a0b80f478f) online.

```

5. Record the WWPN for each FCE-6460-N adapter that is installed on the host. In this example, the WWPN is 200300a0ba0f478F for an adapter on target 0, and would be 200200a0ba0f478F for an adapter on target 3:

```

May 6 10:02:11 solar unix: jnic146x0: Port 011000
(WWN 200200a0b80f478e:200300a0b80f478f) online.
May 6 10:02:11 solar unix: jnic146x3: Port 610300
(WWN 200200a0b80f478e:200200a0b80f478f) online.

```

Attention: In the messages file, notice that each adapter has a network port and a port. Do not use the WWPN for the network port; only use the WWPN for the port.

Opening *jnic146x.conf*: The examples in this procedure assume that you are binding two JNIC FCE-6460 HBAs to the controllers.

Before you begin: See Appendix E, “JNI host bus adapter settings,” on page 111 and modify settings accordingly.

1. Type the following command:

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

2. Look for the section titled Parameters: FcLoopEnabled and FcFabricEnabled. Highlight and copy the following text:


```
-jnic 146x0-FcLoopEnabled=1;
-jnic 146x0-FcFabricEnabled=0;
```

3. Paste the copy immediately following the text that you copied in step 2 on page 80, as shown in the following example:

```
-jnic 146x0-FcLoopEnabled=1;
-jnic 146x0-FcFabricEnabled=0;
-jnic 146x0-FcLoopEnabled=1;
-jnic 146x0-FcFabricEnabled=0;
```

4. Change the target values according to the information that you gathered from the `/var/adm/messages` file. In this example, one controller is on target 0 and the other controller is on target 3:

```
-jnic 146x0-FcLoopEnabled=1;
-jnic 146x0-FcFabricEnabled=0;
-jnic 146x3-FcLoopEnabled=1;
-jnic 146x3-FcFabricEnabled=0;
```

5. Set the `FcLoopEnabled` and `FcFabricEnabled` attributes, as shown in the following example, in which both controllers are attached to a 2 Gb switch:

```
-jnic 146x0-FcLoopEnabled=1;
-jnic 146x0-FcFabricEnabled=1;
-jnic 146x3-FcLoopEnabled=1;
-jnic 146x3-FcFabricEnabled=1;
```

Note: Brocade 2 Gb switches (public):

- `FcLoopEnable=1`
- `FcFabricEnable=1`

Binding the HBAs to the storage controllers: Complete the following steps to set and change the target variables to the specified values:

1. Search for a section titled Configuration Parameters for target to FC device mappings, as shown in the following example:

```
#####
#Configuration Parameters for target to FC device #mappings
#targetX_hba,
#targetX_LunY_hba,
#targetX_wwnn,
#targetX_wpnn,
#targetX_port,

#type: string
#Default
#####
```

2. Ensure that the following lines have been added for each HBA:

```
targetZ_hba = "jnic146x0"
targetZ_wwpn = " "
```

A list is displayed, as shown in the following example:

```
#####
#Configuration Parameters for target to FC device #mappings
#targetX_hba,
#targetX_LunY_hba,
#targetX_wwnn,
#targetX_wpnn,
#targetX_port,

#type: string
#Default

targetZ_hba = 'jnic146x0'
target_Z_wwpn=""
targetY_hba = 'jnic146x0'
target_Y_wwpn=""

#####
```

3. Set the target variables to the values that are specified in the /var/adm/messages file.

The targets in this example are target 0 and target 3:

```
target0_hba = "jnic146x0"
target_Z_wwpn="200300a0b80f478f"
target3_hba="jnic146x3"
target_Y_wwpn=""
```

where Z represents the target number for the first HBA and Y represents the target number for the second HBA from the /var/adm/messages file.

4. Set the target to the WWPN of the first controller, using the value that you recorded from the /var/adm/messages file.

The WWPNs in this example are 200300a0b80f478f for the HBA on target 0 and 200200a0b80f478f for the HBA on target 3:

```
target0_hba = "jnic146x0"
target_0_wwpn="200300a0b80f478f"
target3_hba="jnic146x3"
target_3_wwpn="200200a0b80f478f"
```

5. Type the following command to save the changes to the JNI configuration file:

```
#:wq
```

6. Type the following command to reboot and configure the Solaris host:

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

When you are finished, install the failover driver on the host. See “Installing RDAC on Solaris hosts” for more information.

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 “Installing the JNI adapter driver package” on page 76.

For information about how to set the bindings, see “Modifying the JNI HBA settings” on page 77.

- Check the readme file at the following Web site to ensure that you have the latest version of RDAC:

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

At the Web site, click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for Solaris.

Complete the following steps to install RDAC:

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

Note: 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 “Modifying the JNI HBA settings” on page 77.

Attention: Modifying failover settings in the JNI configuration file after installing RDAC requires the removal of the RDAC from the host. In this case, complete 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, complete 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 (required)" on page 70.

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 FAS*T* 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 that the zones have been enabled.
- Recheck the targets and WWPNs.
- Check to see if the FAS*T* Controllers and expansion drawers are powered on.
- Check to see if the switch or switches 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.

Complete 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 FAS*T* Installation and Support Guide* for the appropriate FAS*T* 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
- Starting the Subsystem Management window
- Renaming storage subsystems
- Performing other storage subsystem management tasks

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.”

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.

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.

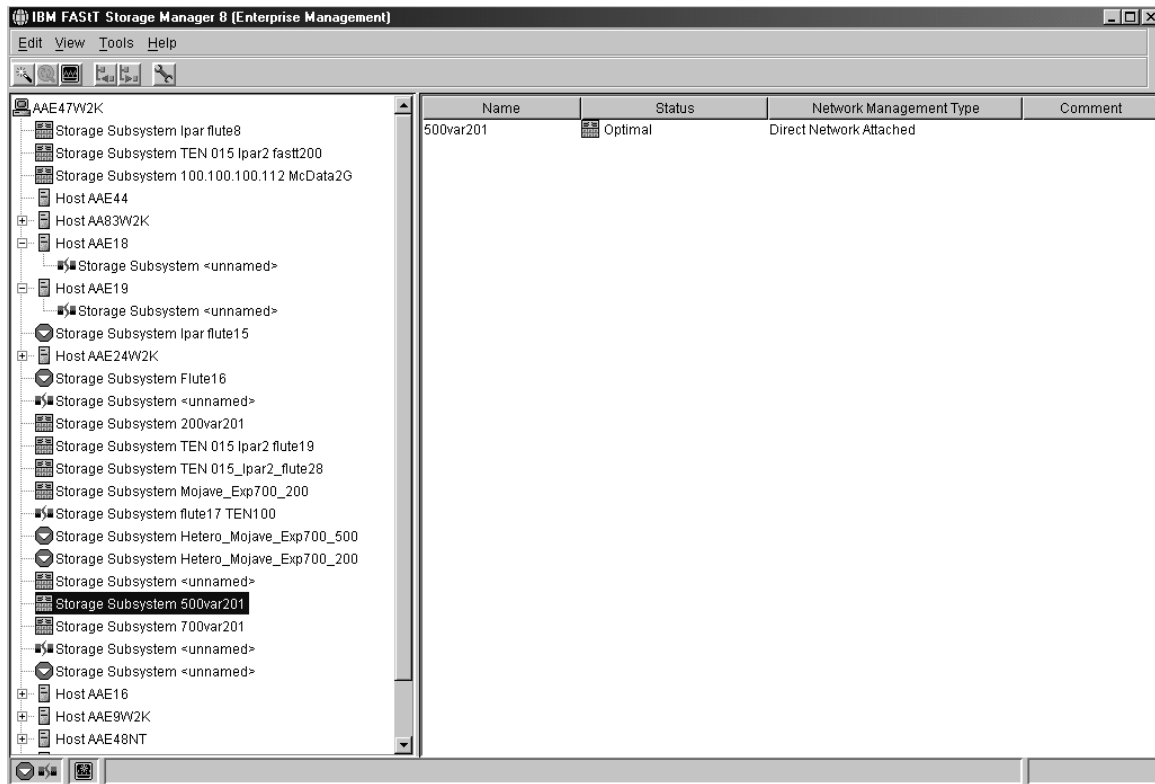
Performing an initial automatic discovery of storage subsystems

Complete the following steps to begin an initial automatic discovery of all the attached hosts and storage subsystems:

1. If you are running the client software from AIX, change to the /SMclient directory and then type `SMclient`. If you are running the client software from Solaris 8, type `SMclient`.
2. Click **Yes** to begin the initial automatic discovery.

Note: 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, as shown in Figure 12 on page 89.



SJ000705

Figure 12. Enterprise Management window

3. 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.
4. 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.
5. Go to “Configuring alert notifications” on page 90.

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 the *IBM FAStT Storage Manager Concepts Guide*, which is available on your installation CD.

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.

Renaming storage subsystems

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

Creating a storage subsystem profile

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.

Complete the following steps to create a storage subsystem profile:

1. Click **View** → **Storage Subsystem Profile** in the Storage Subsystem Management window. The Storage Subsystem Profile window opens.
2. Click **Save As**.
3. Type a file name for the storage subsystem profile and save it to a directory.

If you have a storage server with firmware 04.01.xx.xx or earlier, you must also perform a 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.

Complete the following steps to create a storage subsystem profile by using the Capture State Information feature.

1. Click **Advanced** → **Capture State Information**. The State Capture window opens.
2. Click **Browse** to specify the file name for the state capture information.
3. Click **Start** to start the capture process.
4. Click **Cancel** to close the State Capture window when the capture is complete.

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 17, Chapter 3, “Installing storage management station software on HP-UX systems,” on page 49, or Chapter 4, “Installing storage management station software on Solaris systems,” on page 65 for installation instructions.

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

You can use Storage Manager 8.3 with the following cluster services.

Operating System	High-Availability Cluster Service
AIX	HACMP ES and ESCRM 4.4.1, 4.5 HACMP 5.1 PSSP 3.4 with GPFS 1.5 PSSP 3.5 with GPFS 2.1
HP-UX	MC/Service Guard A.11.09, A.11.13, and A11.14
Solaris	VERITAS Cluster Server 1.3, 2.0, and 3.5 VERITAS Volume Manager 3.2, and 3.5 VERITAS File System 3.2 and 3.5

This document does not describe how to install or configure cluster services. Refer to documentation that is 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.

Using cluster services on HP-UX systems

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.

Using cluster services on Solaris systems

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 just for FAS*T* storage servers, depending on your version of the VERITAS Volume Manager.

For VERITAS Volume Manager 3.1 or earlier, complete the following 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 FAS*T* storage servers only. You do this using option 17 and then option 5 of the submenu of the `vxdiskadm` utility. 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 text editor and make the following changes:

- 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.

Using cluster services on AIX systems

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

General hardware requirements

For the most current information on AIX and device driver versions for FAStT, see the readme file posted at the following Web site:

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

Click the link for your FAStT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAStT Storage Manager for your operating system.

High Availability Cluster Management Program

This section contains general requirements and usage notes for High Availability Cluster Management Program (HACMP) support with Storage Manager 8.3.

Software requirements

The following HACMP software requirements are in addition to the requirements that are specified in Appendix A, "AIX system requirements," on page 103:

- HACMP/ES & ESCRM 4.4.1 with AIX 4.33 or AIX 5.1
- HACMP/ES & ESCRM 4.5 with AIX 5.1 or AIX 5.2
- HACMP 5.1 with AIX 5.1 or AIX 5.2

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.
- The FAStT600 can only be connected to 2 Gb switches, directors, or HBAs. See "FAStT600 restrictions" on page 4 for more information.

Parallel System Support Programs and General Parallel File System

This 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.

Software requirements

The following PSSP and GPFS software requirements are in addition to the requirements that are specified in Appendix A, "AIX system requirements," on page 103.

Table 22. PSSP and GPFS software requirements

PSSP and GPFS versions	Required PSSP APARs	
PSSP 3.4 with GPFS 1.5 PSSP 3.5 with GPFS 2.1	PSSP 3.4: IY31786	PSSP 3.5: IY49062

Configuration limitations

The following limitations apply to PSSP and GPFS configurations:

- Direct connection is not allowed between the host node and a FAStT storage server. Only switched fabric connection is allowed.
- RVSD clusters can support up to two IBM Virtual Shared Disk and RVSD servers for each FAStT partition.
- Single node quorum is not supported in a dual-node GPFS cluster with FAStT disks in the configuration.
- Heterogeneous configurations are not supported.

Other PSSP and GPFS usage notes

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

- Read cache enabled or disabled
- Write cache enabled or disabled
- Cache mirroring enabled or disabled (depending upon the write cache mirroring setting)

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

GPFS, PSSP, and HACMP cluster 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 13 on page 97 shows a cluster configuration that contains a single FAStT storage server, with one to four partitions.

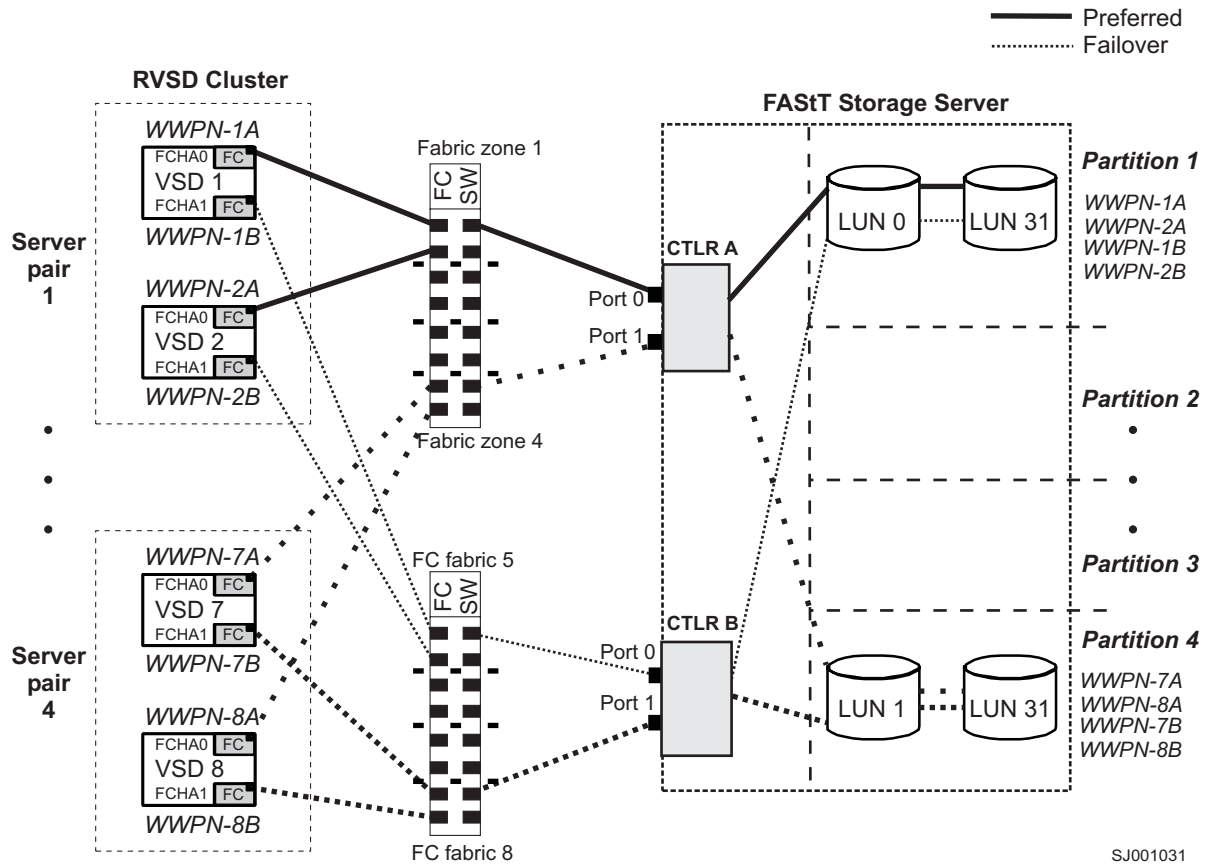


Figure 13. Cluster configuration with single FASTt storage server — one to four partitions

Figure 14 on page 98 shows a cluster configuration that contains three FASTt storage servers, with one partition on each storage server.

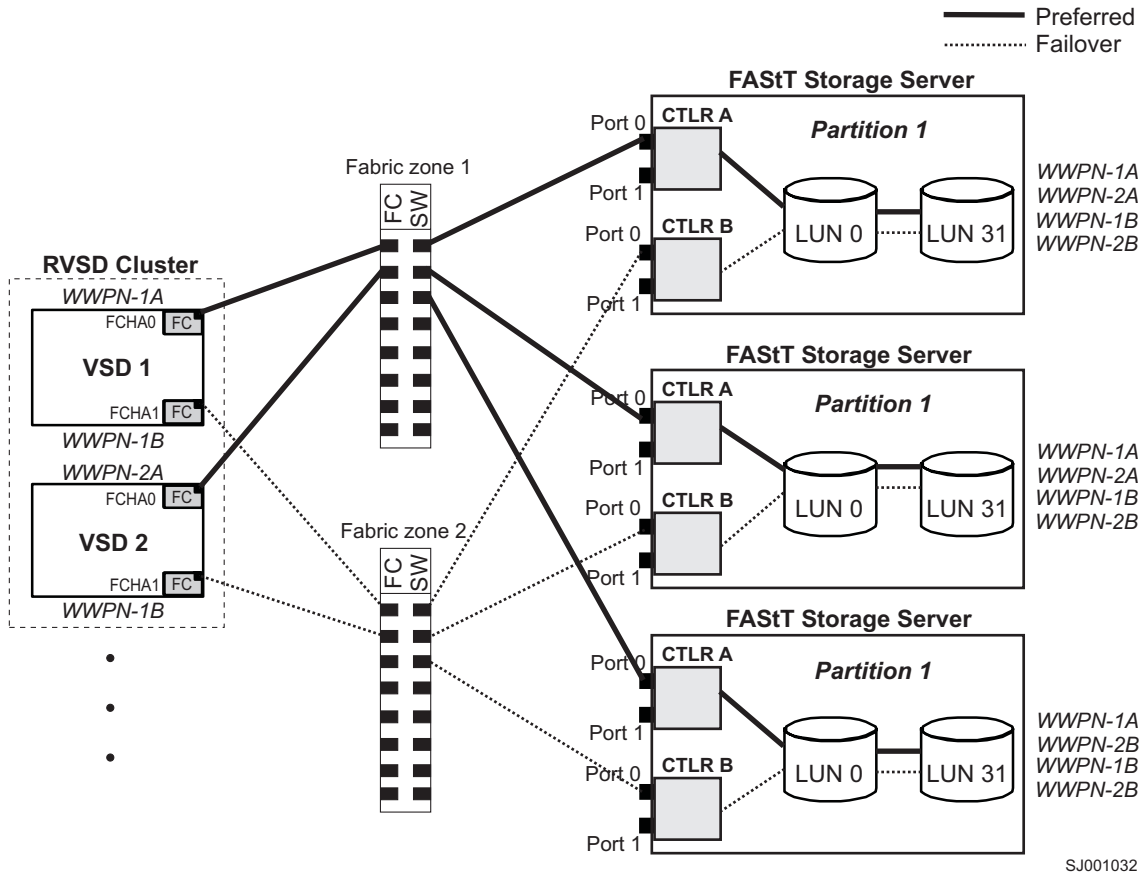


Figure 14. Cluster configuration with three FASTt storage servers — one partition per FASTt

Figure 15 on page 99 shows a cluster configuration that contains four FASTt storage servers, with one partition on each storage server.

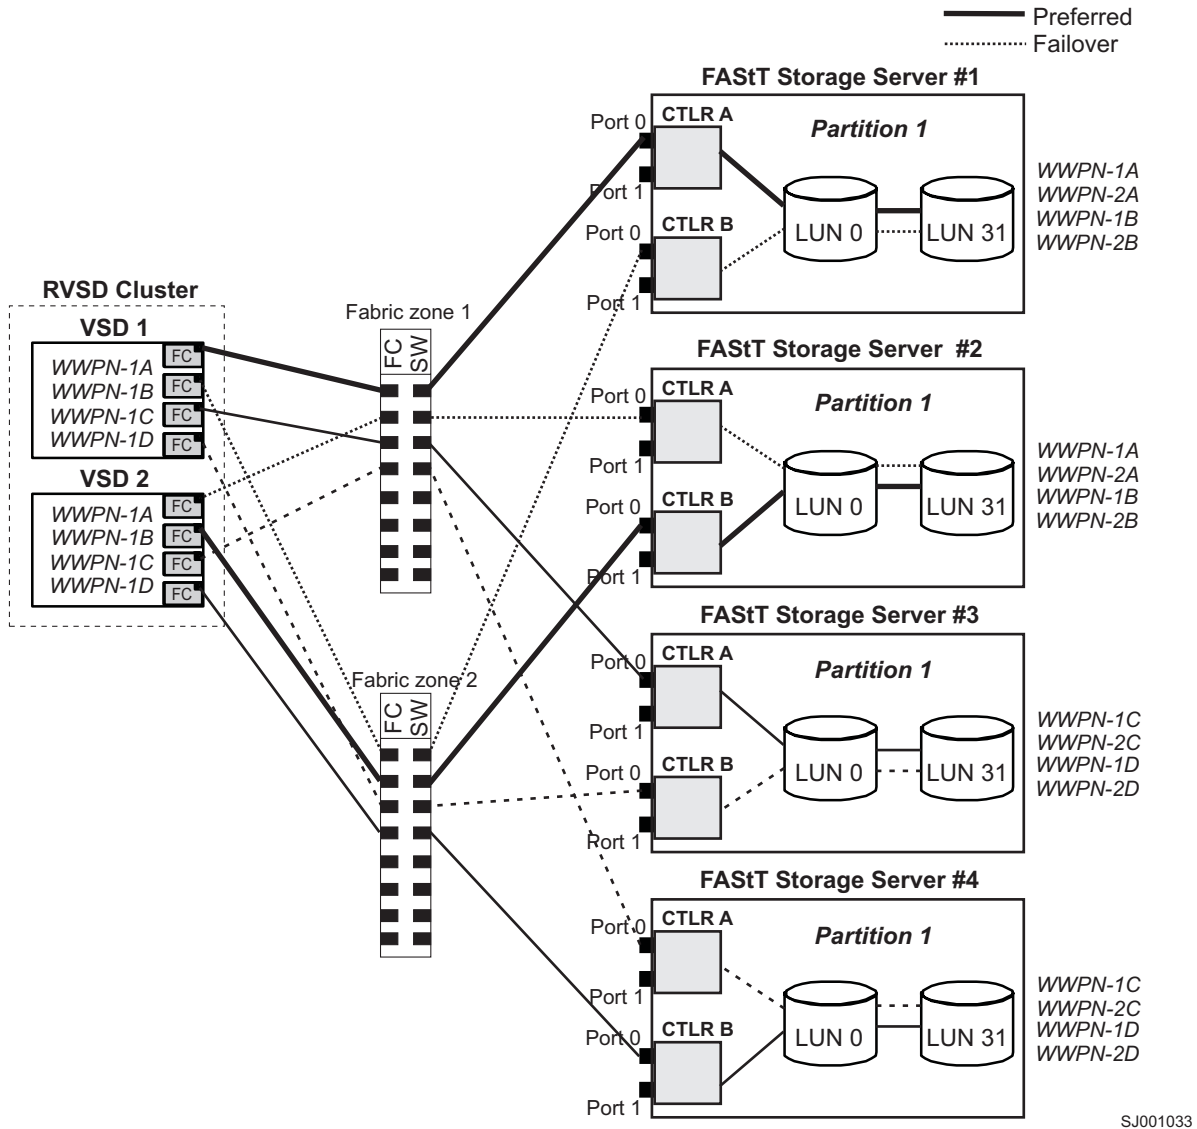


Figure 15. Cluster configuration with four FASTt storage servers — one partition per FASTt

Figure 16 on page 100 shows a cluster configuration that contains two FASTt storage servers, with two partitions on each storage server.

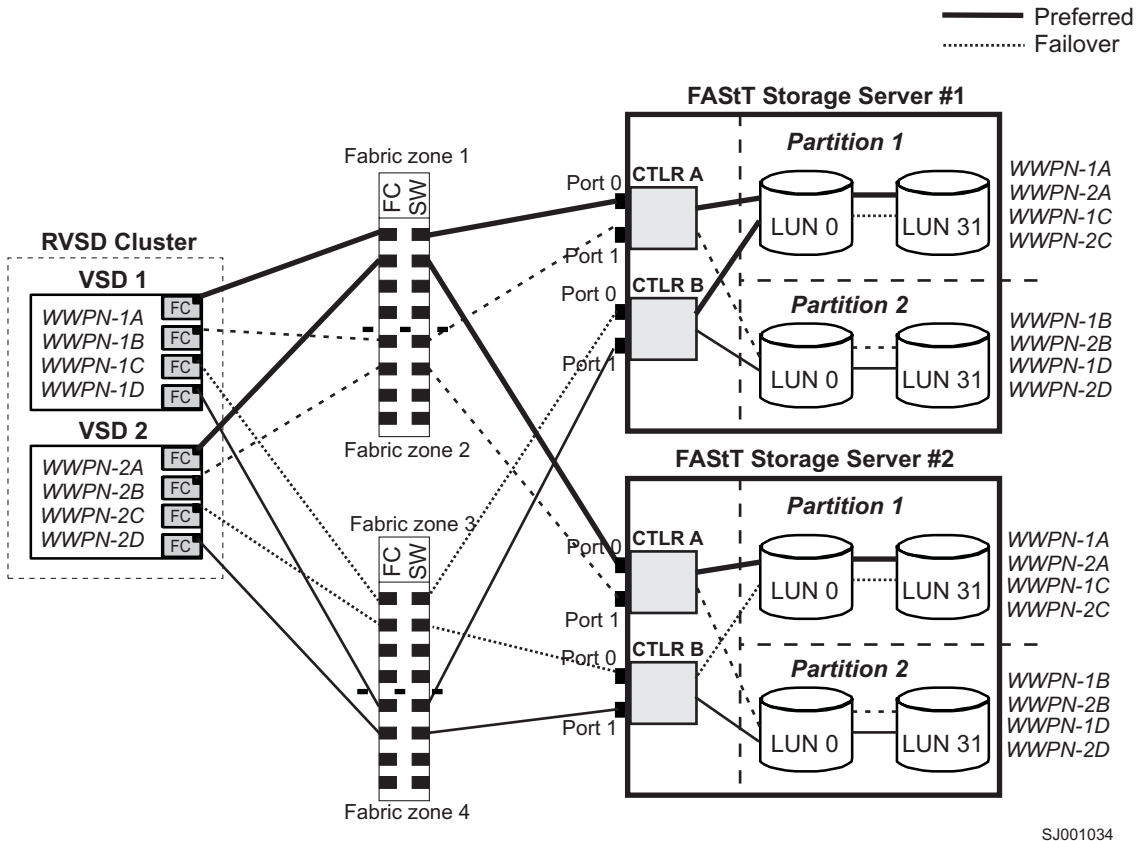


Figure 16. RVSD cluster configuration with two FASTt storage servers — two partitions per FASTt

Figure 17 on page 101 shows an HACMP/GPFS cluster configuration that contains a single FASTt storage server, with one partition.




Figure 17. HACMP/GPFS cluster configuration with one FAST storage server — one partition

Figure 18 on page 102 shows an HACMP/GPFS cluster configuration that contains two FAST storage servers, with two partitions on each storage server.

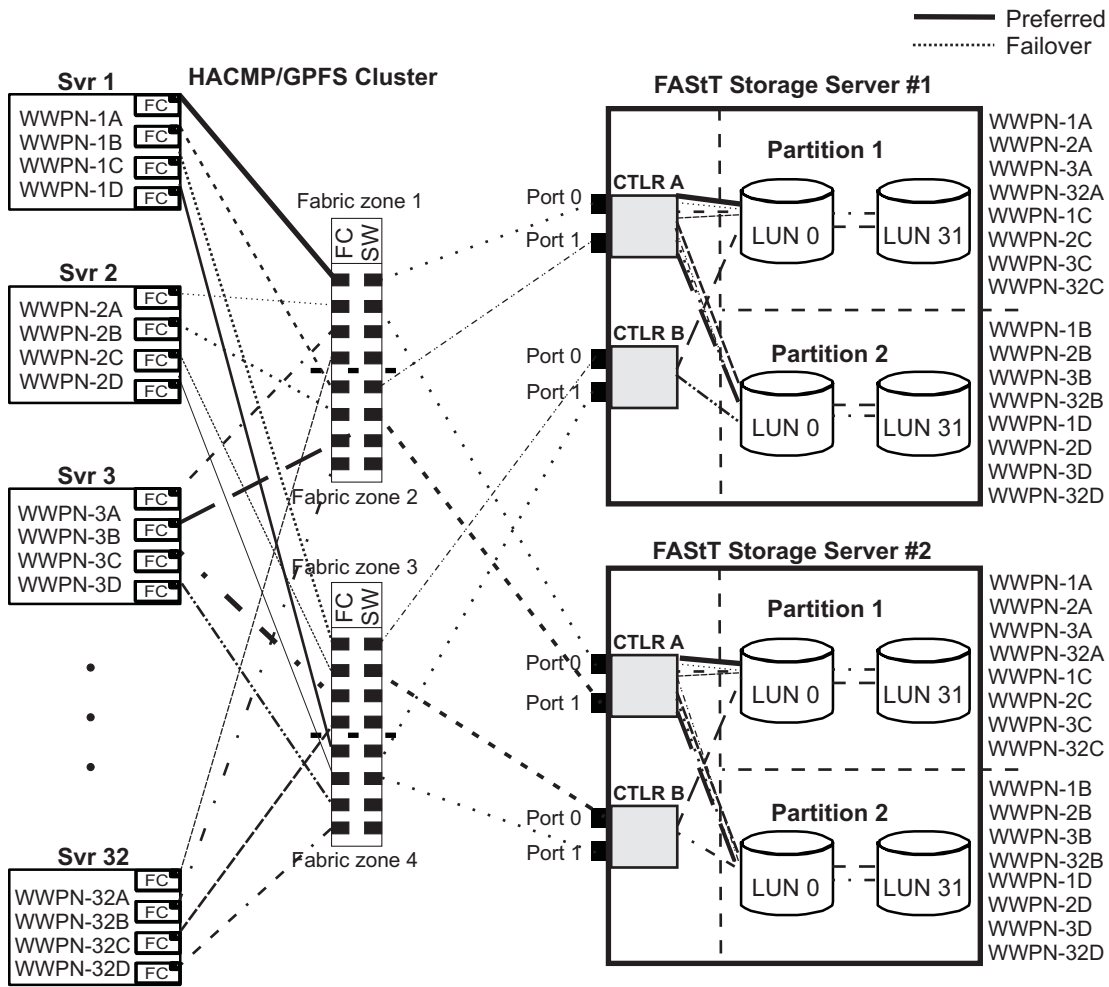


Figure 18. HACMP/GPFS cluster configuration with two FAST storage servers — two partitions per FAST

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 FASTT200, FASTT500, FASTT600, FASTT700, or FASTT900 storage subsystem that is managed by Storage Manager 8.3.

A readme file that contains the latest installation and user information about the storage management software, AIX filesets, and hardware components is located in the AIX FASTT Installation CD. The most recent copy, which supersedes this document, is maintained at the following Web site:

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

Click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for AIX.

Hardware requirements

The following hardware is required at a minimum:

- Any RS/6000® or IBM @server pSeries® server that supports FC 6227, FC 6228 or FC 6239 host bus adapters

Software requirements

The following software must be the specified versions:

- AIX 4.3.3, AIX 5.1 or 5.2
- Filesets associated with IBM RDAC driver installations listed in Table 23, Table 24, or Table 25 on page 104.

RDAC installation requirements

The filesets that are listed in Table 23, Table 24, and Table 25 on page 104 must be the specified versions or later.

Table 23. 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	4.3.3.87
devices.common.IBM.fc.rte	4.3.3.75
devices.pci.df1000f7.com	4.3.3.92
devices.pci.df1000f7.rte	4.3.3.76
devices.pci.df1000f9.rte	4.3.3.76
devices.scsi.sccarray.rte	4.3.3.50

Table 24. Filesets required for AIX 5.1 RDAC

PTF filesets	Version
devices.fcp.disk.array.diag	5.1.0.51
devices.fcp.disk.array.rte	5.1.0.58

Table 24. Filesets required for AIX 5.1 RDAC (continued)

PTF filesets	Version
devices.common.IBM.fc.rte	5.1.0.50
devices.pci.df1000f7.com	5.1.0.57
devices.pci.df1000f7.rte	5.1.0.15
devices.pci.df1000f9.rte	5.1.0.35

Table 25. Filesets required for AIX 5.2 RDAC (with APAR IY35693)

PTF filesets	Version
devices.fcp.disk.array.diag	5.2.0.11
devices.fcp.disk.array.rte	5.2.0.17
devices.common.IBM.fc.rte	5.2.0.10
devices.pci.df1000f7.com	5.2.0.15
devices.pci.df1000f7.rte	5.2.0.10
devices.pci.df1000f9.rte	5.2.0.10

Notes:

1. If you have 2102-F10 storage array devices connected to your system, do not install devices.fcp.disk.array.rte versions 5.1.0.58 or 5.2.0.17 or later. These versions do not support 2102-F10 devices. Contact your IBM technical support representative for assistance.
2. 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/aix.fdc

For more information, see "Installing the RDAC driver" on page 24.

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, FAStT600, FAStT700, or FAStT900 storage subsystem that is managed by Storage Manager 8.3.

A readme 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 supersedes this document, is maintained on the following Web site:

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

Click the link for your FAStT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FAStT Storage Manager for HP-UX.

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
 - PHKL_26519, 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 26 on page 106.

Table 26. 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 a Solaris system must meet to be used in a FASTT200, FASTT500, FASTT600, FASTT700, or FASTT900 storage subsystem that is managed by Storage Manager 8.3.

A readme file that contains the latest installation and user information about the storage management software, and hardware components located in the Solaris FASTT Installation CD. The most recent copy, which supersedes this document, is maintained on the following Web site:

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

Click the link for your FASTT Storage Server. When the page opens, click **Storage Manager v8.3**. Accept the license agreement, then click the **v8.3** link for IBM FASTT Storage Manager for Solaris.

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 6 with the following patches (minimum versions):
 - 105181-33
 - 105356-20
 - 106429-02
- Solaris 7 with the following patches (minimum versions):
 - 106541-23
 - 108376-42 (or later)
 - 107544-03
- Solaris 8 with the following patches (minimum versions):
 - Kernel 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 27. Configuration settings for FCI-1063

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 28. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562

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 28. Configuration settings for FCE-1473/FCE-6460/FCX2-6562/FCC2-6562 (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/genjnicnf` reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicnf
```

Table 29. Configuration settings for FCE-1063/FCE2-1063/FCE-6410/FCE2-6410

Original value	New value	Comment
FcLoopEnabled = 1	FcLoopEnabled = 0	Uncomment the line.
FcFabricEnabled = 0	FcFabricEnabled = 1	Uncomment the line.
FailoverDelay = 30;	FailoverDelay = 60;	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/genjnicnf` reconfigure script from the Solaris shell.

```
# /etc/raid/bin/genjnicnf
```

- 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 30. Configuration settings for FC64-1063

Original value	New value	Comments
fca_nport = 0;	fca_nport = 1;	

Table 30. Configuration settings for FC64-1063 (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/genscsiconf` reconfigure script from the shell prompt.

```
# /etc/raid/bin/genscsiconf
```

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
@server
FlashCopy
Netfinity

pSeries
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 peripheral component interconnect (PCI) bus. AGP reduces the overall cost of creating high-end graphics subsystems by using existing system memory.

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

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

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

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

AGP. See *accelerated graphics port*.

AL_PA. See *arbitrated loop physical address*.

arbitrated loop. One of three existing fibre-channel topologies, in which 2 - 126 ports are interconnected serially in a single loop circuit. Access to the Fibre Channel-Arbitrated Loop (FC-AL) is controlled by an arbitration scheme. The FC-AL topology supports all classes of service and guarantees in-order delivery of FC frames when the originator and responder are on the same FC-AL. The default topology for the disk array is arbitrated loop. An arbitrated loop is sometimes referred to as a Stealth Mode.

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

array. A set of fibre-channel hard drives that are logically grouped together to form a redundant array of independent disks (RAID) set.

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

ATA. See *AT-attached*.

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

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

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

AWT. See *Abstract Windowing Toolkit*.

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

BIOS. See *basic input/output system*.

BOOTP. See *bootstrap protocol*.

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

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

bridge group. A bridge and the collection of devices connected to it.

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

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

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

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

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

CRC. See *cyclic redundancy check*.

CRT. See *cathode ray tube*.

CRU. See *customer replaceable unit*.

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

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

dac. See *disk array controller*.

dar. See *disk array router*.

DASD. See *direct access storage device*.

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

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

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

DHCP. See *Dynamic Host Configuration Protocol*.

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

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

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

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

DMA. See *direct memory access*.

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

DRAM. See *dynamic random access memory*.

Dynamic Host Configuration Protocol (DHCP). A protocol defined by the Internet Engineering Task Force that is used for dynamically assigning Internet Protocol (IP) addresses to computers in a network.

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

ECC. See *error correction coding*.

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

EISA. See *Extended Industry Standard Architecture*.

electrically erasable programmable read-only memory (EEPROM). A type of memory chip which can retain its contents without consistent electrical power. Unlike the PROM which can be programmed only once, the EEPROM can be erased electrically. Because it can only be reprogrammed a limited number of times before it wears out, it is appropriate for storing small amounts of data that are changed infrequently.

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

environmental services monitor (ESM) canister. A component in a drive enclosure that monitors the environmental condition of the components in that enclosure. Not all storage subsystems have ESM canisters.

E_port. See *expansion port*.

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

ESD. See *electrostatic discharge*.

ESM canister. See *environmental services monitor canister*.

EXP. See *expansion unit*.

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

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

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

fabric. A Fibre Channel entity which interconnects and facilitates logins of N_ports attached to it. The fabric is responsible for routing frames between source and destination N_ports using address information in the frame header. A fabric can be as simple as a point-to-point channel between two N-ports, or as complex as a frame-routing switch that provides multiple and redundant internal pathways within the fabric between F_ports.

fabric port (F_port). In a fabric, an access point for connecting a user's N_port. An F_port facilitates N_port

logins to the fabric from nodes connected to the fabric. An F_port is addressable by the N_port connected to it. See also *fabric*.

FAStT MSJ. See *FAStT Management Suite Java*.

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

FC. See *fibre channel*.

FC-AL. See *arbitrated loop*.

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

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

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

Fibre Channel Protocol (FCP) for small computer system interface (SCSI). A high-level fibre-channel mapping layer (FC-4) that uses lower-level fibre-channel (FC-PH) services to transmit SCSI commands, data, and status information between a SCSI initiator and a SCSI target across the FC link by using FC frame and sequence formats.

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

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

F_port. See *fabric port*.

FRU. See *field replaceable unit*.

GBIC. See *gigabit interface converter*

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

graphical user interface (GUI). A type of computer interface that presents a visual metaphor of a real-world scene, often of a desktop, by combining high-resolution

graphics, pointing devices, menu bars and other menus, overlapping windows, icons, and the object-action relationship.

GUI. See *graphical user interface*.

HBA. See *host bus adapter*.

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

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

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

host computer. See *host*.

host group. An entity in the storage partition topology that defines a logical collection of host computers that require shared access to one or more logical drives.

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

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

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

IBMSAN driver. The device driver that is used in a Novell NetWare environment to provide multipath input/output (I/O) support to the storage controller.

IC. See *integrated circuit*.

IDE. See *integrated drive electronics*.

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

Industry Standard Architecture (ISA). Unofficial name for the bus architecture of the IBM PC/XT personal computer. This bus design included expansion slots for plugging in various adapter boards. Early versions had an 8-bit data path, later expanded to 16 bits. The "Extended Industry Standard Architecture" (EISA) further expanded the data path to 32 bits. See also *Extended Industry Standard Architecture*.

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

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

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

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

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

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

IP. See *Internet Protocol*.

IPL. See *initial program load*.

IRQ. See *interrupt request*.

ISA. See *Industry Standard Architecture*.

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

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

JRE. See *Java Runtime Environment*.

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

LAN. See *local area network*.

LBA. See *logical block address*.

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

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

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

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

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

loop port. A node port (N_port) or fabric port (F_port) that supports arbitrated loop functions associated with an arbitrated loop topology.

LUN. See *logical unit number*.

MAC. See *medium access control*.

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

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

MCA. See *micro channel architecture*.

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

MIB. See *management information base*.

micro channel architecture (MCA). Hardware that is used for PS/2 Model 50 computers and above to

provide better growth potential and performance characteristics when compared with the original personal computer design.

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

network management station (NMS). In the Simple Network Management Protocol (SNMP), a station that runs management application programs that monitor and control network elements.

NMI. See *non-maskable interrupt*.

NMS. See *network management station*.

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

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

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

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

N_port. See *node port*.

NVS. See *nonvolatile storage*.

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

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

ODM. See *Object Data Manager*.

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

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

PDF. See *portable document format*.

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

peripheral component interconnect local bus (PCI local bus). A local bus for PCs, from Intel, that provides a high-speed data path between the CPU and up to 10 peripherals (video, disk, network, and so on).

The PCI bus coexists in the PC with the Industry Standard Architecture (ISA) or Extended Industry Standard Architecture (EISA) bus. ISA and EISA boards plug into an IA or EISA slot, while high-speed PCI controllers plug into a PCI slot. See also *Industry Standard Architecture*, *Extended Industry Standard Architecture*.

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

port. A part of the system unit or remote controller to which cables for external devices (such as display stations, terminals, printers, switches, or external storage units) are attached. The port is an access point for data entry or exit. A device can contain one or more ports.

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

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

private loop. A freestanding arbitrated loop with no fabric attachment. See also *arbitrated loop*.

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

PTF. See *program temporary fix*.

RAID. See *redundant array of independent disks*.

RAM. See *random-access memory*.

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

RDAC. See *redundant disk array controller*.

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

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

redundant array of independent disks (RAID). A collection of disk drives that appears as a single volume to the server and are fault tolerant through mirroring or parity checking.

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.

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*.

SATA. See *serial ATA*.

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

SCSI. See *small computer system interface*.

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

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

serial ATA. The standard for a high-speed alternative to small computer system interface (SCSI) hard drives. The SATA-1 standard is equivalent in performance to a 10 000 RPM SCSI drive.

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

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

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

SFP. See *small form-factor pluggable*.

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

SL_port. See *segmented loop port*.

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

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

SMruntime. A Java compiler for the SMclient.

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

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

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

SNMP. See *Simple Network Management Protocol* and *SNMPv1*.

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

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

SRAM. See *static random access memory*.

SSA. See *serial storage architecture*.

static random access memory (SRAM). Random access memory based on the logic circuit known as flip-flop. It is called static because it retains a value as long as power is supplied, unlike dynamic random access memory (DRAM), which must be regularly refreshed. It is however, still volatile, meaning that it can lose its contents when the power is turned off.

storage area network (SAN). A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services. See also *fabric*.

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

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

storage partition topology. In the FAStT Storage Manager client, the Topology view of the Mappings window displays the default host group, the defined host group, the host computer, and host-port nodes. The host port, host computer, and host group topological elements must be defined to grant access to host computers and host groups using logical drive-to-LUN mappings.

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

sweep method. A method of sending Simple Network Management Protocol (SNMP) requests for information to all the devices on a subnet by sending the request to every device in the network.

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

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

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

TL_port. See *translated loop port*.

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

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

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

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

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

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

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

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

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

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

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

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

write-once read many (WORM). Any type of storage medium to which data can be written only a single time, but can be read from any number of times. After the data is recorded, it cannot be altered.

WWN. See *worldwide name*.

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

Index

A

- about this document xi
- access volume 2
- access volumes 9
- address
 - for e-mail comments xxii
 - for reader comment form xxii
 - of IBM xxii
 - of the IBM director of licensing 115
- agent software
 - package 6
- AIX
 - client software 19
 - configuration of storage subsystems 20, 25
 - direct-attached configuration 23
 - firmware
 - requirements 17
 - upgrading 21
 - hardware requirements 103
 - identifying
 - controller ID numbers 33
 - device names and bus numbers 34
 - FlashCopy volume types 42
 - logical drives by operating system device names 40
 - NVSRAM
 - requirements 17
 - upgrading 22
 - operating system requirements 13, 14, 17
 - RDAC driver
 - configuring 25
 - installing 24, 25
 - requirements 103
 - requirements
 - client software 19
 - firmware 17
 - hardware 17, 103
 - operating system 13, 14
 - RDAC driver 103
 - software 103
 - system 103
 - restrictions 9, 18
 - SAN-attached configuration 24
 - SMclient 20
 - SMruntime 20
 - software requirements 103
 - system requirements 103
 - volumes, redistributing in case of failure 47
- alert notifications, configuring 90
- attributes
 - dac 27, 31
 - dar 26, 30
 - definitions 26
 - hdisk 28, 31, 32
 - LUN 28, 31, 32
 - RDAC driver 25
- audience of this document xi

- Auto Volume Transfer (AVT) 63, 85

B

- bus numbers and device names, identifying 34

C

- cache mirroring 47
- capture state information feature 90
- catastrophic failure 90
- changing the host type 87
- client software
 - AIX 19, 20
 - HP-UX 51
 - installing 50
 - package 5, 6
 - Solaris 67, 68
 - uninstalling 71
- cluster services
 - AIX requirements 95
 - HACMP ES and ESCRM 93
 - hardware requirements 93
 - HP-UX requirements 93
 - MC/Service Guard 93
 - PSSP with GPFS 93
 - Solaris requirements 93, 94
 - system dependencies 93
 - VERITAS Cluster Server 93
 - VERITAS Volume Manager 93, 94
- commands
 - fget_config 36
 - lsattr 39
 - lsdev 35
- comments about this document, how to send xxii
- configuration details, MC Service Guard 109
- configuration recovery 90
- configuring
 - alert notifications 90
 - direct-attached 23, 57, 74
 - RDAC driver 25, 94
 - SAN-attached 24, 57, 75
 - Solaris with JNI host bus adapter cards 76, 77
 - storage subsystems
 - AIX 20, 23, 24
 - HP-UX 54, 57
 - procedures 88
 - Solaris 71, 74, 75
- controller state capture 90
- controllers
 - firmware 1, 54, 72
 - ID numbers, identifying 33
 - IP addresses 14

D

- dac attributes 27, 31

- dar
 - two dars on one subsystem, causes 26
- dar attributes 26, 30
- DCE (dynamic capacity expansion) 44
- default partitioning, Solaris devices 74
- delayed_fail 43
- devices
 - identification 33
 - names and bus numbers, identifying 34
 - Solaris, default partitioning for 74
- direct (out-of-band) management method 9
- direct-attached configuration 23, 57, 74
- disk configuration information 90
- documentation
 - documents 1
 - FAST xii
 - FAST Storage Manager Version 8.3 xiii
 - FAST200 xviii
 - FAST500 xvii
 - FAST700 xvi
 - FAST900 xv
 - files 1
 - online help xx, 2, 89, 91
 - related xix
 - Web sites xxi
- DVE (dynamic volume expansion) 44
- dynamic capacity expansion (DCE) 44
- dynamic tracking 44
- dynamic volume expansion (DVE) 44

E

- e-mail address xxii, 90
- edition notice ii
- enabling multipath I/O with PV-links 57
- Enterprise Management window
 - description 6
 - online help 2

F

- failure support
 - cluster services 93
 - RDAC driver 6
 - redistributing volumes 47, 63, 85
- fast I/O failure 43
- fast_fail 43
- FAST documentation xii
- FAST Storage Manager
 - related documents xix
 - Version 8.3 library xiii
- FAST, introduction to 14
- FAST200 Storage Server library xviii
- FAST500 Storage Server library xvii
- FAST600
 - restrictions 4, 23, 24, 75, 109
- FAST600 FC Storage Server library xiv
- FAST600 Turbo feature 3, 5
- FAST700 Storage Server library xvi
- FAST900 FC Storage Server library xv
- fcsci attributes 43, 44

- fget_config command 36
- filesets 103
- firmware
 - AIX
 - requirements 17
 - upgrading 21
 - description 1
 - downloading latest versions 21, 72
 - HP-UX
 - requirements 49
 - upgrading 55
 - Solaris
 - requirements 65
 - upgrading 72

H

- HACMP ES and ESCRM 93
- hardware requirements
 - AIX 17, 103
 - HP-UX 49, 105
 - Solaris 65, 107
- HBA
 - binding to storage controllers 77, 78, 80, 81
 - modifying settings 77, 79
 - selecting host port
 - AIX 22
 - HP-UX 55
 - Solaris 73
- hdisk attributes 28, 31, 32
- help xx, 2, 89, 91
- heterogeneous environment 88
- high-availability cluster services 93
- host group, setting up
 - AIX 22
 - HP-UX 55
 - Solaris 73
- host software
 - HP-UX 52, 53
 - packages 6
 - Solaris 69, 70, 71
- host type 87
- host-agent (in-band) management method 7
- how this document is organized xx
- how to send your comments xxii
- HP-UX
 - client software 50, 51
 - cluster services, requirements 93
 - configuration of storage subsystems 54
 - direct-attached configuration 57
 - firmware
 - requirements 49
 - upgrading 55
 - hardware requirements 49, 105
 - host software
 - installing 52
 - requirements 52
 - NVSRAM
 - requirements 49
 - upgrading 55
 - operating system requirements 13, 14

HP-UX (*continued*)

- requirements
 - client software 51
 - cluster services 93
 - firmware 49
 - hardware 49, 105
 - host software 52
 - NVSRAM 49
 - operating system 13, 14
 - software 105
 - system 105
- SAN-attached configuration 57
- SMagent 52
- SMclient 51
- SMruntime 51
- SMutil 53
- software requirements 105
- system requirements 105
- volumes, redistributing in case of failure 63

HyperTerminal version 14

I

IBM

- address xxii
- director of licensing address 115
- e-mail address xxii, 90

identifying

- AIX
 - controller ID numbers 33
 - device names and bus numbers 34
 - logical drives by operating system device names 40

in-band (host-agent) management method 7

installing

- AIX
 - client software 19, 20
 - RDAC driver 24
- client software
 - AIX 19
 - HP-UX 50
 - SMclient 20, 51, 68
 - SMruntime 20, 51, 67
 - Solaris 67
- host software
 - HP-UX 52
 - SMagent 52, 70
 - SMutil 53, 70
 - Solaris 69
- HP-UX
 - client software 50, 51
 - host software 52
 - SMutil 53
- post-installation tasks 87
- process diagram for 11
- sequence of 6
- Solaris
 - client software 67, 68
 - host software 69, 70
 - RDAC driver 83

IP addresses for FASiT controllers 14

J

JNI

- cards 76, 77
- modifying HBA settings 77, 79
- settings 111

L

- loop settings 77, 79
- lsattr command 39
- lsdev command 35
- LUN attributes 31
- LUNs
 - and access volumes 9
 - attributes 28, 31, 32

M

management method

- direct (out-of-band) 9
- host-agent (in-band) 7

management station 1

- See storage management station

MC Service Guard 109

MC/Service Guard 93

multipath I/O with PV-links 57

N

notice, edition ii

notices

- general 115

NVSRAM

- AIX
 - requirements 17, 18, 50
 - downloading latest versions 21, 72
- HP-UX
 - requirements 49
 - upgrading 55
- Solaris
 - requirements 65
 - upgrading 72

O

online help xx, 2, 89, 91

operating system

- AIX 13, 14, 17, 103
- HP-UX 13, 14, 105
 - requirements 49
- requirements
 - AIX 17
 - HP-UX 49
 - Solaris 65
- Solaris 13, 14, 107
 - requirements 65

organization of this document xx

out-of-band (direct) management method 9

overview of heterogeneous hosts 88

P

- partitioning for Solaris devices (default) 74
- premium features 2
 - copy services 2
 - enabling 4
 - storage partitioning 3
 - host group 22, 55, 73
- prerequisites
 - AIX
 - client software 19
 - firmware 17
 - hardware 17, 103
 - operating system 13, 14
 - RDAC driver 103
 - software 103
 - system 103
 - cluster services 93
 - HP-UX
 - client software 51
 - cluster services 93
 - firmware 49
 - hardware 49, 105
 - host software 52
 - NVSRAM 49
 - operating system 13, 14
 - software 105
 - system 105
 - Solaris
 - client software 67
 - cluster services 93, 94
 - firmware 65
 - hardware 65, 107
 - host software 69
 - NVSRAM 65
 - operating system 13, 14
 - RDAC driver 69
 - software 107
 - system 107
- products, developed 115
- PSSP with GPFS 93
- PV-links 57

R

- RDAC driver
 - AIX
 - configuring 25
 - installing 24, 25
 - requirements 103
 - attributes 25
 - description 6
 - IDs 94
 - installing 82
 - Solaris
 - installing 82, 83
 - requirements 69
 - uninstalling 71
- reader comment form xxii
- readme file 1
- recover configuration 90

- redistributing volumes in case of failure 47, 63, 85
- renaming storage subsystem 90
- requirements
 - AIX
 - client software 19
 - firmware 17
 - hardware 17, 103
 - operating system 13, 14
 - RDAC driver 103
 - software 103
 - system 103
 - cluster services 93
 - HP-UX
 - client software 51
 - cluster services 93
 - firmware 49
 - hardware 49, 105
 - host software 52
 - NVSRAM 49
 - operating system 13, 14
 - software 105
 - system 105
 - Solaris
 - client software 67
 - cluster services 93, 94
 - firmware 65
 - hardware 65, 107
 - host software 69
 - NVSRAM 65
 - operating system 13, 14
 - RDAC driver 69
 - software 107
 - system 107
- resources
 - documents 1
 - e-mail address xxii, 90
 - files 1
 - online help xx, 2, 89, 91
 - Web sites xxi, 4
- restrictions
 - AIX 9, 18
 - FASTT600 4, 23, 24, 75, 109
 - in-band management 7
- runtime software
 - package 5

S

- SAN-attached configuration 24, 57, 75
- sending your comments to IBM xxii
- services offered in the U.S.A. 115
- setting
 - IP addresses 14
 - NMS for SNMP notification 90
- SMagent
 - description 6
 - HP-UX installing 52
 - installing 70
 - Solaris installing 70
- SMclient
 - AIX installing 20

- SMclient (*continued*)
 - description 5
 - HP-UX installing 51
 - Solaris installing 68
 - SMruntime
 - AIX installing 20
 - description 5
 - HP-UX installing 51
 - Solaris installing 67
 - SMutil
 - description 6
 - HP-UX installing 53
 - Solaris installing 70
 - SNMP notification 90
 - software package
 - RDAC 6
 - SMagent 6
 - SMclient 5
 - SMruntime 5
 - SMutil 6
 - software requirements
 - AIX 103
 - HP-UX 105
 - Solaris 107
 - Solaris
 - client software 67, 71
 - cluster services requirements 93, 94
 - configuration of storage subsystem 71
 - devices, default partitioning 74
 - direct-attached configuration 74
 - firmware
 - requirements 65
 - upgrading 72
 - hardware requirements 65, 107
 - host software 69, 71
 - JNI host bus adapter cards 76, 77
 - NVSRAM
 - requirements 65
 - upgrading 72
 - operating system requirements 13, 14
 - RDAC driver
 - installing 83
 - requirements 69
 - uninstalling 71
 - requirements
 - client software 67
 - cluster services 93
 - firmware 65
 - hardware 65, 107
 - host software 69
 - NVSRAM 65
 - operating system 13, 14
 - RDAC driver 69
 - software 107
 - system 107
 - SAN-attached configuration 75
 - SMagent
 - installing 70
 - SMclient
 - installing 68
 - Solaris (*continued*)
 - SMruntime
 - installing 67
 - SMutil
 - installing 70
 - software requirements 107
 - system requirements 107
 - volumes, redistributing in case of failure 85
 - storage management station 1
 - Storage Manager software
 - installation sequence 6
 - introduction 1
 - software packages 5
 - uninstalling, Solaris 71
 - where to obtain 1
 - storage subsystem
 - cluster services 93
 - configuring
 - AIX 20
 - HP-UX 54
 - procedures 88
 - Solaris 71
 - introduction 1
 - logical drives
 - identifying by operating system device names 40
 - management methods
 - direct (out-of-band) 9
 - host-agent (in-band) 7
 - management tasks 91
 - renaming 90
 - storage subsystem profile, saving 90
 - Subsystem Management window
 - description 6
 - online help 2
 - switch
 - in a SAN-attached configuration 24, 57, 75
 - modifying JNI HBA settings 77, 79
 - zoning 24, 57, 75
 - system requirements
 - AIX 103
 - HP-UX 105
 - Solaris 107
- T**
- trademarks 115
- U**
- uninstalling Storage Manager software, Solaris 71
 - universal transport mechanism 25
 - upgrading
 - AIX
 - NVSRAM 22
 - HP-UX
 - firmware 55
 - NVSRAM 55
 - Solaris
 - firmware 72
 - NVSRAM 72

- user tasks
 - FAST600 Fibre Channel Storage Server xiv
- utility software
 - package 6
- utm device 25

V

- variable settings for JNI adapter cards 77
- VERITAS Cluster Server 93
- VERITAS Volume Manager 93, 94
- volumes
 - redistributing in case of failure 47, 63, 85

W

- Web sites xxi, 4, 18, 24, 50, 95
- who should
 - read this document xi
 - use this document xi

Z

- zoning switches 24, 57, 75

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. GC26-7521-03

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

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

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

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

Please tell us how we can improve this book:

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

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
Dept. FTFA/Building 050-1
5600 Cottle Road
San Jose, CA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GC26-7521-03



Spine information:



IBM TotalStorage FAS*StT* Storage
Manager Version 8.3

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